

HIGH-RESOLUTION FINITE VOLUME MODELING OF WAVE PROPAGATION IN ORTHOTROPIC POROELASTIC MEDIA

GRADY I. LEMOINE^{*†}, M. YVONNE OU[†], AND RANDALL J. LEVEQUE^{*}

Abstract. Poroelasticity theory models the dynamics of porous, fluid-saturated media. It was pioneered by Maurice Biot in the 1930s through 1960s, and has applications in several fields, including geophysics and modeling of *in vivo* bone. A wide variety of methods have been used to model poroelasticity, including finite difference, finite element, pseudospectral, and discontinuous Galerkin methods. In this work we use a Cartesian-grid high-resolution finite volume method to numerically solve Biot's equations in the time domain for orthotropic materials, with the stiff relaxation source term in the equations incorporated using operator splitting. This class of finite volume method has several useful properties, including the ability to use wave limiters to reduce numerical artifacts in the solution, ease of incorporating material inhomogeneities, low memory overhead, and an explicit time-stepping approach. To the authors' knowledge, this is the first use of high-resolution finite volume methods to model poroelasticity. The solution code uses the CLAWPACK finite volume method software, which also includes block-structured adaptive mesh refinement in its AMRCLAW variant. We present convergence results for known analytic plane wave solutions, achieving second-order convergence rates outside of the stiff regime of the system. Our convergence rates are degraded in the stiff regime, but we still achieve similar levels of error on the finest grids examined. We also demonstrate good agreement against other numerical results from the literature. To aid in reproducibility, we provide all of the code used to generate the results of this paper, at https://bitbucket.org/grady_lemoine/poro-2d-cartesian-archive.

Key words. poroelastic, wave propagation, finite-volume, high-resolution, running head, stiff relaxation, operator splitting

1. Introduction. Poroelasticity theory is a homogenized model for solid porous media containing fluids that can flow through the pore structure. This field was pioneered by Maurice A. Biot, who developed his theory of poroelasticity from the 1930s through the 1960s; a summary of much of Biot's work can be found in his 1956 and 1962 papers [4, 5, 6]. Biot theory uses linear elasticity to describe the solid portion of the medium (often termed the *skeleton* or *matrix*), linearized compressible fluid dynamics to describe the fluid portion, and Darcy's law to model the aggregate motion of the fluid through the matrix. While it was originally developed to model fluid-saturated rock and soil, Biot theory has also been used in underwater acoustics [10, 30, 31], and to describe wave propagation in *in vivo* bone [19, 20, 29].

Biot theory predicts rich and complex wave phenomena within poroelastic materials. Three different types of waves appear: fast P waves analogous to standard elastic P waves, in which the fluid and matrix show little relative motion, and typically compress or expand in phase with each other; shear waves analogous to elastic S waves; and slow P waves, where the fluid expands while the solid contracts, or vice versa. The slow P waves exhibit substantial relative motion between the solid and fluid compared to waves of the other two types. The viscosity of the fluid dissipates poroelastic waves as they propagate through the medium, with the fast P and S waves being lightly damped and the slow P wave strongly damped. The viscous dissipation also causes slight dispersion in the fast P and S waves, and strong dispersion in the slow P wave.

A variety of different numerical approaches have been used to model poroelasticity. Carcione, Morency, and Santos provide a thorough review of the previous literature [14]. The earliest numerical work in poroelasticity seems to be that of Garg [27], using a finite difference method in 1D. Finite difference and pseudospectral methods have continued to be popular since then, with further work by Mikhailenko [39], Hasanzadeh [33], Dai et al. [22], and more recently Chiavassa and Lombard [17], among others. Finite element approaches began being used in the 1980s, with Santos and Oreña's work [43] being one of the first. Boundary element methods have also been used, such as in the work of Attenborough, Berry, and Chen [2]. Spectral element methods have also been used in both the frequency domain [24] and the time domain [40]. With the recent rise of discontinuous Galerkin methods, DG has been applied to poroelasticity in several works, such as that of de la Puente et al. [23]. There have also been semi-analytical approaches to solving the poroelasticity equations, such as that of Detournay and Cheng [26], who analytically obtain a solution in the Laplace transform domain, but are forced to use an approximate inversion procedure to return to the

^{*}Department of Applied Mathematics, University of Washington, Guggenheim Hall Box 352420, Seattle, WA 98195. Supported in part by NIH grant 5R01AR53652-2 and NSF grant DMS-0914942

[†]Department of Mathematical Sciences, University of Delaware, 501 Ewing Hall, Newark, DE 19716. Supported in part by NSF-DMS Mathematical Biology Grant 0920852 and NIH CBER PILOT Grant 322 159

[‡]Corresponding author, email: g1@uw.edu

time domain. Finally, there has been significant work on inverse problems in poroelasticity, for which various forward solvers have been used; of particular note is the paper of Buchanan, Gilbert, and Khashanah [9], who used the finite element method (specifically the FEMLAB software package) to obtain time-harmonic solutions for cancellous bone as part of an inversion scheme to estimate poroelastic material parameters, and the later papers of Buchanan and Gilbert [7, 8], where the authors instead used numerical contour integration of the Green’s function. Numerical work in the 1970s and 1980s focused on isotropic poroelasticity, with the earliest work on anisotropic poroelasticity being by Carcione in 1996 [12].

A major theme in time-domain numerical modeling of poroelasticity has been the difficulty of handling the viscous dissipation term, which has its own intrinsic time scale and causes the poroelasticity system to be stiff, at least if low-frequency waves are being considered. (The time scales associated with dissipation are independent of frequency, so at higher frequencies there is less separation between them and the time scale associated with wave motion — in other words, the system is not stiff if sufficiently high-frequency waves are being considered.) This viscous dissipation is particularly problematic for the slow P wave. While it can still be addressed if Biot’s equations are solved as a unified system (e.g. with an implicit time-integration method), since the viscous dissipation term is easy to solve analytically in isolation, operator splitting approaches have also been a popular way to address this issue; we adopt this approach as well. Carcione and Quiroga-Goode used operator splitting in conjunction with a pseudospectral method [15], and Chiavassa and Lombard use it with a finite difference approach [17]. De la Puente et al. also investigate an operator splitting approach, and encounter difficulties in obtaining a fast rate of convergence due to the stiffness of the dissipation term, which pushed the problems they investigated toward the diffusive limit of the Biot system [23].

Our work in this paper solves a velocity-stress formulation of Biot linear orthotropic poroelasticity theory using Cartesian-grid high-resolution finite volume methods. These methods are memory-efficient explicit techniques designed to model hyperbolic systems, and can include wave limiters designed to reduce the effect of numerical artifacts in the solution. Since they are based on solving Riemann problems at grid interfaces, it is also straightforward to include material inhomogeneities between cells using these methods. To our knowledge, this is the first use of finite volume methods to model poroelasticity. We employ the CLAWPACK finite volume method package [45], which offers the use of operator splitting to model viscous dissipation, and includes optional Berger-Colella-Oliger block-structured adaptive mesh refinement [3] to improve solution efficiency for larger problems. We verify our code both against known analytical solutions and against other numerical solutions from the poroelasticity literature.

2. Poroelasticity theory for transversely isotropic materials. Biot’s equations of poroelasticity are a complicated system of PDEs that exhibit rich and varied behaviors. The reader is encouraged to refer to a detailed treatment of the subject, such as chapter 7 of Carcione’s book [13], but we also provide an overview here. After a review of the basic relations modeling the behavior of a poroelastic medium in sections 2.1 through 2.4, section 2.5 presents the linear first-order system of PDEs that forms the basis of our numerical model. Section 2.6 then defines an energy norm for our state vector that solves some of the scaling issues associated with modeling poroelasticity in SI units. The matrix created for this energy norm allows us to explore some useful properties of the model: section 2.7 provides a concise proof that our first-order system is hyperbolic, and section 2.8 exhibits a strictly convex entropy function for the system, which has implications for the correctness of our numerical solutions that we explore later in section 3.3.

2.1. Stress-strain relation. We assume the constituent material of the solid matrix is isotropic and the anisotropy of the solid matrix results purely from the microstructure. We also assume that the anisotropy has a specific form — that the medium is orthotropic, possessing three orthogonal planes of symmetry, and is isotropic with respect to some axis of symmetry. This type of anisotropy is common in engineering composites [28], and in biological materials [21], as well as being present in certain types of stone. Let the z -axis be this axis of symmetry. The elastic stiffness tensor \mathbf{C} of such an orthotropic, transversely isotropic medium contains five independent components. In its principal axes, and using shorthand notation, \mathbf{C} can

be arranged as

$$\mathbf{C} = \begin{pmatrix} c_{11} & c_{12} & c_{13} & 0 & 0 & 0 \\ c_{12} & c_{11} & c_{13} & 0 & 0 & 0 \\ c_{13} & c_{13} & c_{33} & 0 & 0 & 0 \\ 0 & 0 & 0 & c_{55} & 0 & 0 \\ 0 & 0 & 0 & 0 & c_{55} & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{c_{11}-c_{12}}{2} \end{pmatrix}, \quad (2.1)$$

with the stress tensor and engineering strains arranged into vectors $\boldsymbol{\tau}$ and \mathbf{e} of the form

$$\boldsymbol{\tau} = (\tau_{11} \ \tau_{22} \ \tau_{33} \ \tau_{23} \ \tau_{13} \ \tau_{12})^T \quad \text{and} \quad \mathbf{e} = (\epsilon_{11} \ \epsilon_{22} \ \epsilon_{33} \ 2\epsilon_{23} \ 2\epsilon_{13} \ 2\epsilon_{12})^T. \quad (2.2)$$

Note the factor of 2 applied to the shear strains to convert from the tensor strain $\epsilon_{ij} = \frac{1}{2}(\partial_i u_j + \partial_j u_i)$ to the engineering strain. With this shorthand notation, the stress-strain relation is

$$\boldsymbol{\tau} = \mathbf{C}\mathbf{e}. \quad (2.3)$$

2.2. Energy densities and the dissipation potential. One useful property of the poroelasticity system is that it admits an energy density, from which we can define an energy norm that we will use extensively.

This section is mainly based on Biot's 1956 papers [4, 5] and chapter 7 of Carcione's book [13]. All formulations are in terms of the following variables:

1. \mathbf{u} , the displacement vector of the solid matrix
2. $\mathbf{w} := \phi(\mathbf{U} - \mathbf{u})$, the relative motion of the fluid scaled by the porosity, where \mathbf{U} is the displacement vector of the pore fluid and ϕ is the porosity of the medium
3. $\zeta := -\nabla \cdot \mathbf{w}$, the variation in fluid content

2.2.1. Strain energy. In terms of the undrained elasticity tensor \mathbf{C}^u and the strain components of the solid matrix $e_{ii}^{(m)} := \partial_i u_i$ and $e_{ij}^{(m)} := \partial_i u_j + \partial_j u_i$, $i \neq j$, the strain energy density of the Biot model for 3D transversely isotropic materials, with z -axis being the axis of symmetry, is given by

$$2V = c_{11}^u \left(e_{11}^{(m)2} + e_{22}^{(m)2} \right) + c_{33}^u e_{33}^{(m)2} + 2c_{12}^u e_{11}^{(m)} e_{22}^{(m)} + 2c_{13}^u \left(e_{11}^{(m)} + e_{22}^{(m)} \right) e_{33}^{(m)} \\ + c_{55}^u (e_{23}^{(m)2} + e_{13}^{(m)2}) + c_{66}^u e_{12}^{(m)2} - 2\alpha_1 M \left(e_{11}^{(m)} + e_{22}^{(m)} \right) \zeta - 2\alpha_3 M e_{33}^{(m)} \zeta + M \zeta^2, \quad (2.4)$$

where $c_{66}^u = \frac{c_{11}^u - c_{12}^u}{2}$ for a transversely isotropic material and the undrained elastic constants c_{ij}^u are related to those of the dry matrix, $c_{ij}^{(m)}$, via

$$c_{ij}^u = c_{ij}^{(m)} + M\alpha_i\alpha_j, \quad i, j = 1, \dots, 6 \quad (2.5)$$

$$\boldsymbol{\alpha} := (\alpha_1, \alpha_1, \alpha_3, 0, 0, 0) \quad (2.6)$$

$$\alpha_1 := 1 - \frac{c_{11}^{(m)} + c_{12}^{(m)} + c_{13}^{(m)}}{3K_s} \quad (2.7)$$

$$\alpha_3 := 1 - \frac{2c_{13}^{(m)} + c_{33}^{(m)}}{3K_s} \quad (2.8)$$

$$M := \frac{K_s^2}{K_s [1 + \phi(K_s/K_f - 1)] - \left(2c_{11}^{(m)} + c_{33}^{(m)} + 2c_{12}^{(m)} + 4c_{13}^{(m)} \right) / 9}. \quad (2.9)$$

Here K_s and K_f are the bulk moduli of the constituent material of the solid matrix and of the pore fluid, respectively. The total stress (solid matrix plus pore pressure) acting on a volume element of the medium are given by the derivatives of strain energy with respect to the associated strains,

$$\tau_{ij} = \frac{\partial V}{\partial e_{ij}^{(m)}}, \quad (2.10)$$

or, in short notation,

$$\tau_I = \sum_{J=1}^6 c_{IJ}^u e_J^{(m)} - M \alpha_I \zeta. \quad (2.11)$$

Similarly, the pore pressure p is

$$p = \frac{\partial V}{\partial \zeta} = M \left(\zeta - \sum_{j=1}^3 \alpha_j e_{jj}^{(m)} \right). \quad (2.12)$$

Most of the work in this paper is for 2D plane-strain conditions in the x - z plane. For these conditions, we will later need a function \tilde{V} that maps from the pressure and the in-plane stress components τ_{11} , τ_{13} , τ_{33} at a point to the strain energy density at that point. To obtain this, we will first derive the strains as a function of stress for plane-strain conditions by setting $e_{12}^{(m)} = e_{22}^{(m)} = e_{23}^{(m)} = 0$ in equations (2.11) and (2.12), then solving for the remaining strains and the variation in fluid content:

$$\mathbf{e}_{\text{plane}} = \mathbf{S}_{\text{plane}} \boldsymbol{\tau}_{\text{plane}}, \quad (2.13)$$

where the matrices $\mathbf{e}_{\text{plane}}$, $\mathbf{S}_{\text{plane}}$, and $\boldsymbol{\tau}_{\text{plane}}$ are

$$\mathbf{e}_{\text{plane}} = \begin{pmatrix} e_{11}^{(m)} & e_{33}^{(m)} & e_{13}^{(m)} & \zeta \end{pmatrix}^T, \quad \boldsymbol{\tau}_{\text{plane}} = \begin{pmatrix} \tau_{11} & \tau_{33} & \tau_{13} & p \end{pmatrix}^T, \quad (2.14)$$

and

$$\mathbf{S}_{\text{plane}} = \begin{pmatrix} \frac{c_{33}^{(m)}}{c_{11}^{(m)} c_{33}^{(m)} - (c_{13}^{(m)})^2} & -\frac{c_{13}^{(m)}}{c_{11}^{(m)} c_{33}^{(m)} - (c_{13}^{(m)})^2} & 0 & \frac{\alpha_1 c_{33}^{(m)} - \alpha_3 c_{13}^{(m)}}{c_{11}^{(m)} c_{33}^{(m)} - (c_{13}^{(m)})^2} \\ -\frac{c_{13}^{(m)}}{c_{11}^{(m)} c_{33}^{(m)} - (c_{13}^{(m)})^2} & \frac{c_{11}^{(m)}}{c_{11}^{(m)} c_{33}^{(m)} - (c_{13}^{(m)})^2} & 0 & \frac{\alpha_3 c_{11}^{(m)} - \alpha_1 c_{13}^{(m)}}{c_{11}^{(m)} c_{33}^{(m)} - (c_{13}^{(m)})^2} \\ 0 & 0 & \frac{1}{c_{55}^{(m)}} & 0 \\ \frac{\alpha_1 c_{33}^{(m)} - \alpha_3 c_{13}^{(m)}}{c_{11}^{(m)} c_{33}^{(m)} - (c_{13}^{(m)})^2} & \frac{\alpha_3 c_{11}^{(m)} - \alpha_1 c_{13}^{(m)}}{c_{11}^{(m)} c_{33}^{(m)} - (c_{13}^{(m)})^2} & 0 & \frac{1}{M} + \frac{\alpha_1^2 c_{33}^{(m)} + \alpha_3^2 c_{11}^{(m)} - 2\alpha_1 \alpha_3 c_{13}^{(m)}}{c_{11}^{(m)} c_{33}^{(m)} - (c_{13}^{(m)})^2} \end{pmatrix}. \quad (2.15)$$

Substituting (2.13) through (2.15) into (2.4) yields the strain energy as a function of stress for plane strain conditions, which can be expressed as the quadratic form

$$\tilde{V} := \frac{1}{2} \boldsymbol{\tau}_{\text{plane}}^T \mathbf{S}_{\text{plane}} \boldsymbol{\tau}_{\text{plane}} \quad (2.16)$$

Note the use of the drained elastic coefficients ($c_{11}^{(m)}$, etc.) rather than the undrained coefficients (c_{11}^u). We expect this quadratic form to be positive-definite on physical grounds — if it were not, then it would be possible to deform the medium or change its fluid content without doing work.

2.2.2. Kinetic energy and the dissipation potential. For anisotropic poroelastic media, the kinetic energy density has the form

$$T = \frac{1}{2} \left(\dot{\mathbf{u}}^T \mathbf{P} \dot{\mathbf{u}} + 2 \dot{\mathbf{u}}^T \mathbf{R} \dot{\mathbf{U}} + \dot{\mathbf{U}}^T \mathbf{T} \dot{\mathbf{U}} \right), \quad (2.17)$$

where \mathbf{R} is the induced mass matrix. Assume all the three matrices can be diagonalized in the same coordinate system so that $\mathbf{P} = \text{diag}(a_1, a_2, a_3)$, $\mathbf{R} = \text{diag}(r_1, r_2, r_3)$ and $\mathbf{T} = \text{diag}(t_1, t_2, t_3)$. By looking at the special case of no relative motion between fluid and solid, it can be shown that the mass coefficients satisfy the two equations (given as (7.169) in Carcione [13])

$$a_i + r_i = (1 - \phi) \rho_s, \quad r_i + t_i = \phi \rho_f. \quad (2.18)$$

Defining the velocity variables

$$\mathbf{v} := \dot{\mathbf{u}}, \quad \mathbf{q} := \dot{\mathbf{w}}, \quad (2.19)$$

the kinetic energy density can be expressed as

$$\begin{aligned} T &= \frac{1}{2} \sum_{i=1}^3 \left[(1-\phi) \rho_s v_i^2 - r_i (v_i - \dot{U}_i)^2 + \phi \rho_f \dot{U}_i^2 \right] \\ &= \frac{1}{2} \sum_{i=1}^3 \left[\rho v_i^2 + 2\rho_f q_i v_i + \left(\frac{\rho_f \phi - r_i}{\phi^2} \right) q_i^2 \right] =: \tilde{T}(\mathbf{v}, \mathbf{q}), \end{aligned} \quad (2.20)$$

where ρ_s and ρ_f are the constituent solid density and pore fluid density, respectively and $\rho = (1-\phi)\rho_s + \phi\rho_f$ is the bulk density of the medium. The induced mass parameters r_i are related to the tortuosity T_i by

$$r_i = \phi \rho_f (1 - T_i). \quad (2.21)$$

Assuming the pore fluid flow is of the Poiseuille type, the dissipation potential Φ_D in an anisotropic medium in terms of the dynamic viscosity η of pore fluid and the permeability tensor \mathbf{K} is

$$\Phi_D = \frac{1}{2} \sum_{i,j=1}^3 \eta (\mathbf{K}^{-1})_{ij} q_i q_j. \quad (2.22)$$

Assume that \mathbf{K} has the same principal directions as \mathbf{P} , \mathbf{R} and \mathbf{T} with eigenvalues $\kappa_1, \kappa_2, \kappa_3$. Then we have

$$\Phi_D = \frac{1}{2} \sum_{i=1}^3 \frac{\phi^2 \eta}{\kappa_i} (\dot{U}_i - v_i)^2 = \frac{1}{2} \sum_{i=1}^3 \frac{\eta}{\kappa_i} q_i^2 =: \tilde{\Phi}_D(\mathbf{q}) \quad (2.23)$$

2.3. Equations of motion. The equations of motion for the solid part are given in terms of the energy densities and dissipation potential as

$$\partial_t \left(\frac{\partial T}{\partial v_i} \right) + \frac{\partial \Phi_D}{\partial v_i} = \sum_{j=1}^3 \partial_j \left(\frac{\partial V}{\partial e_{ij}^{(m)}} + \phi p \delta_{ij} \right), \quad (2.24)$$

or equivalently

$$\partial_t \left(\frac{\partial \tilde{T}}{\partial v_i} - \phi \frac{\partial \tilde{T}}{\partial q_i} \right) - \phi \frac{\partial \tilde{\Phi}_D}{\partial q_i} = \sum_{j=1}^3 \partial_j \left(\frac{\partial V}{\partial e_{ij}^{(m)}} + \phi p \delta_{ij} \right). \quad (2.25)$$

Similarly, the equation of motion for the fluid part is

$$\partial_t \left(\frac{\partial T}{\partial \dot{U}_i} \right) + \frac{\partial \Phi_D}{\partial \dot{U}_i} = -\phi \sum_{j=1}^3 \partial_j \left(\frac{\partial V}{\partial \zeta} \delta_{ij} \right), \quad (2.26)$$

or equivalently

$$\phi \partial_t \left(\frac{\partial \tilde{T}}{\partial q_i} \right) + \phi \frac{\partial \tilde{\Phi}_D}{\partial q_i} = -\phi \sum_{j=1}^3 \partial_j \left(\frac{\partial V}{\partial \zeta} \delta_{ij} \right), \quad (2.27)$$

which reduces to

$$-\partial_i p = \rho_f \dot{v}_i + \left(\frac{\rho_f \phi - r_i}{\phi^2} \right) \dot{q}_i + \frac{\eta}{\kappa_i} q_i, \quad i = 1, 2, 3. \quad (2.28)$$

The equations of motion for the fluid-solid composite are obtained by adding (2.25) with (2.27), giving

$$\sum_{j=1}^3 \partial_j \tau_{ij} = \partial_t \left(\frac{\partial \tilde{T}}{\partial v_i} \right) = \rho \dot{v}_i + \rho_f \dot{q}_i. \quad (2.29)$$

2.4. Governing equations for plane-strain case. Since the material is assumed isotropic in the x - y plane, we consider the plane strain problem in the x - z plane. The governing equations are obtained by suppressing the y -component (subscript 2) of \mathbf{U} , \mathbf{u} , \mathbf{w} , \mathbf{q} and those terms of $e_{ij}^{(m)}$ with $i = 2$ or $j = 2$ in V , T , \tilde{T} , Φ_D , and $\tilde{\Phi}_D$. (While nonzero out-of-plane stresses do arise in a plane-strain problem, they do not produce in-plane motion, and can be ignored for purposes of studying the in-plane dynamics of the medium.) We obtain two types of governing equation:

- Stress-strain relations, obtained by differentiating (2.12) and (2.11) for $I = 1, 3, 5$ with respect to time,

$$\partial_t \tau_{xx} = c_{11}^u \partial_x v_x + c_{13}^u \partial_z v_z + \alpha_1 M (\partial_x q_x + \partial_z q_z) + \partial_t s_1 \quad (2.30)$$

$$\partial_t \tau_{zz} = c_{13}^u \partial_x v_x + c_{33}^u \partial_z v_z + \alpha_3 M (\partial_x q_x + \partial_z q_z) + \partial_t s_3 \quad (2.31)$$

$$\partial_t \tau_{xz} = c_{55}^u (\partial_z v_x + \partial_x v_z) + \partial_t s_5 \quad (2.32)$$

$$\partial_t p = -\alpha_1 M \partial_x v_x - \alpha_3 M \partial_z v_z - M (\partial_x q_x + \partial_z q_z) + \partial_t s_f, \quad (2.33)$$

where s_1 , s_3 , s_5 , and s_f are the solid and fluid external sources.

- Equations of motion

$$\rho \partial_t v_x + \rho_f \partial_t q_x = \partial_x \tau_{xx} + \partial_z \tau_{xz} \quad (2.34)$$

$$\rho \partial_t v_z + \rho_f \partial_t q_z = \partial_x \tau_{xz} + \partial_z \tau_{zz} \quad (2.35)$$

$$\rho_f \partial_t v_x + m_1 \partial_t q_x + \left(\frac{\eta}{\kappa_1} \right) q_x = -\partial_x p \quad (2.36)$$

$$\rho_f \partial_t v_z + m_3 \partial_t q_z + \left(\frac{\eta}{\kappa_3} \right) q_z = -\partial_z p, \quad (2.37)$$

where $m_i := \frac{\rho_f \phi - r_i}{\phi^2} = \frac{\rho_f T_i}{\phi}$.

2.5. Governing equations as a linear first-order system. Solving (2.34) and (2.36) for $\partial_t v_x$ and $\partial_t q_x$, we obtain

$$\partial_t v_x = \frac{1}{\Delta_1} \left(m_1 \partial_x \tau_{xx} + m_1 \partial_z \tau_{xz} + \rho_f \partial_x p + \rho_f \frac{\eta}{\kappa_1} q_x \right) \quad (2.38)$$

$$\partial_t q_x = \frac{1}{\Delta_1} \left(-\rho_f \partial_x \tau_{xx} - \rho_f \partial_z \tau_{xz} - \rho \partial_x p - \rho \frac{\eta}{\kappa_1} q_x \right), \quad (2.39)$$

where $\Delta_1 := \rho m_1 - \rho_f^2$. Similarly, (2.35) and (2.37) lead to

$$\partial_t v_z = \frac{1}{\Delta_3} \left(m_3 \partial_x \tau_{xz} + m_3 \partial_z \tau_{zz} + \rho_f \partial_z p + \rho_f \frac{\eta}{\kappa_3} q_z \right) \quad (2.40)$$

$$\partial_t q_z = \frac{1}{\Delta_3} \left(-\rho_f \partial_x \tau_{xz} - \rho_f \partial_z \tau_{zz} - \rho \partial_z p - \rho \frac{\eta}{\kappa_3} q_z \right), \quad (2.41)$$

where $\Delta_3 := \rho m_3 - \rho_f^2$. Combining the stress-strain relations (2.30)-(2.33) with equations (2.38)-(2.41), we obtain the 8×8 linear first-order system

$$\partial_t \mathbf{Q} + \mathbf{A} \partial_x \mathbf{Q} + \mathbf{B} \partial_z \mathbf{Q} = \mathbf{D} \mathbf{Q} + \partial_t \mathbf{s}, \quad (2.42)$$

where

$$\mathbf{Q} = (\tau_{xx} \quad \tau_{zz} \quad \tau_{xz} \quad v_x \quad v_z \quad p \quad q_x \quad q_z)^T \quad (2.43)$$

$$\mathbf{A} = - \begin{pmatrix} 0 & 0 & 0 & c_{11}^u & 0 & 0 & \alpha_1 M & 0 \\ 0 & 0 & 0 & c_{13}^u & 0 & 0 & \alpha_3 M & 0 \\ 0 & 0 & 0 & 0 & c_{55}^u & 0 & 0 & 0 \\ \frac{m_1}{\Delta_1} & 0 & 0 & 0 & 0 & \frac{\rho_f}{\Delta_1} & 0 & 0 \\ 0 & 0 & \frac{m_3}{\Delta_3} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -\alpha_1 M & 0 & 0 & -M & 0 \\ -\frac{\rho_f}{\Delta_1} & 0 & 0 & 0 & 0 & -\frac{\rho}{\Delta_1} & 0 & 0 \\ 0 & 0 & -\frac{\rho_f}{\Delta_3} & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \quad (2.44)$$

$$\mathbf{B} = - \begin{pmatrix} 0 & 0 & 0 & 0 & c_{13}^u & 0 & 0 & \alpha_1 M \\ 0 & 0 & 0 & 0 & c_{33}^u & 0 & 0 & \alpha_3 M \\ 0 & 0 & 0 & c_{55}^u & 0 & 0 & 0 & 0 \\ 0 & 0 & \frac{m_1}{\Delta_1} & 0 & 0 & 0 & 0 & 0 \\ 0 & \frac{m_3}{\Delta_3} & 0 & 0 & 0 & \frac{\rho_f}{\Delta_3} & 0 & 0 \\ 0 & 0 & 0 & 0 & -\alpha_3 M & 0 & 0 & -M \\ 0 & 0 & -\frac{\rho_f}{\Delta_1} & 0 & 0 & 0 & 0 & 0 \\ 0 & -\frac{\rho_f}{\Delta_3} & 0 & 0 & 0 & -\frac{\rho}{\Delta_3} & 0 & 0 \end{pmatrix} \quad (2.45)$$

$$\mathbf{D} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \frac{\rho_f \eta}{\Delta_1 \kappa_1} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & \frac{\rho_f \eta}{\Delta_3 \kappa_3} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -\frac{\rho \eta}{\Delta_1 \kappa_1} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -\frac{\rho \eta}{\Delta_3 \kappa_3} \end{pmatrix}, \quad (2.46)$$

and

$$\mathbf{s} = (s_1 \quad s_3 \quad s_5 \quad 0 \quad 0 \quad s_f \quad 0 \quad 0)^T. \quad (2.47)$$

It is this system (2.42) that forms the basis for our numerical work. Note that while the coefficient matrices \mathbf{A} , \mathbf{B} , and \mathbf{D} are defined in the material principal axes, we can extend this system to model media where the principal axes are different from the global x - z axes through an appropriate transformation of the state variables in \mathbf{Q} that come from vector and tensor quantities, and application of the chain rule in the partial derivatives with respect to the spatial variables. In such cases, we refer to the principal directions as the 1 and 3 axes to distinguish them from the computational x and z axes.

It is worth noting that (2.42) is not just a generic “black box” equation, but one of a very specific type: a first-order hyperbolic system with a stiff relaxation source term. We will prove hyperbolicity in section 2.7, and the source term $\mathbf{D}\mathbf{Q}$ shows itself to be of relaxation type by having only zero and negative eigenvalues — in the absence of the spatial derivative terms, it would cause the solution to decay exponentially toward the null space $\mathcal{N}(\mathbf{D})$, and even with the other terms present we can expect it to keep the solution close to $\mathcal{N}(\mathbf{D})$. (Whether the relaxation term really is stiff depends on the other time scales of the particular problem being solved, but it is stiff for some of the problems considered here.) On the subject of time scales, because \mathbf{D} is extremely sparse, we can immediately read off the eigenvalues associated with dissipation in the 1 and 3 axes — respectively, $-\frac{\rho \eta}{\Delta_1 \kappa_1}$ and $-\frac{\rho \eta}{\Delta_3 \kappa_3}$ — and so define the characteristic time for decay in each axis as the negative inverse of these eigenvalues,

$$\tau_{d1} := \frac{\Delta_1 \kappa_1}{\rho \eta}, \quad \tau_{d3} := \frac{\Delta_3 \kappa_3}{\rho \eta}. \quad (2.48)$$

2.6. Energy norm. Let $\mathcal{E} := \tilde{T} + \tilde{V}$ be the total mechanical energy per unit volume in a representative element, where \tilde{T} is the kinetic energy from (2.20), and \tilde{V} is the strain energy from (2.16). For the subsequent analysis, we will use its Hessian with respect to the state variables in \mathbf{Q} , which is the symmetric matrix

$$\mathbf{E} = \begin{pmatrix} \frac{c_{33}^{(m)}}{c_{11}^{(m)} c_{33}^{(m)} - (c_{13}^{(m)})^2} & -\frac{c_{13}^{(m)}}{c_{11}^{(m)} c_{33}^{(m)} - (c_{13}^{(m)})^2} & 0 & 0 & 0 & \frac{\alpha_1 c_{33}^{(m)} - \alpha_3 c_{13}^{(m)}}{c_{11}^{(m)} c_{33}^{(m)} - (c_{13}^{(m)})^2} & 0 & 0 \\ -\frac{c_{13}^{(m)}}{c_{11}^{(m)} c_{33}^{(m)} - (c_{13}^{(m)})^2} & \frac{c_{11}^{(m)}}{c_{11}^{(m)} c_{33}^{(m)} - (c_{13}^{(m)})^2} & 0 & 0 & 0 & \frac{\alpha_3 c_{11}^{(m)} - \alpha_1 c_{13}^{(m)}}{c_{11}^{(m)} c_{33}^{(m)} - (c_{13}^{(m)})^2} & 0 & 0 \\ 0 & 0 & \frac{1}{c_{55}^{(m)}} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \rho & 0 & 0 & \rho_f & 0 \\ 0 & 0 & 0 & 0 & \rho & 0 & 0 & \rho_f \\ \frac{\alpha_1 c_{33}^{(m)} - \alpha_3 c_{13}^{(m)}}{c_{11}^{(m)} c_{33}^{(m)} - (c_{13}^{(m)})^2} & \frac{\alpha_3 c_{11}^{(m)} - \alpha_1 c_{13}^{(m)}}{c_{11}^{(m)} c_{33}^{(m)} - (c_{13}^{(m)})^2} & 0 & 0 & 0 & \frac{1}{M} + \frac{\alpha_1^2 c_{33}^{(m)} + \alpha_3^2 c_{11}^{(m)} - 2\alpha_1 \alpha_3 c_{13}^{(m)}}{c_{11}^{(m)} c_{33}^{(m)} - (c_{13}^{(m)})^2} & 0 & 0 \\ 0 & 0 & 0 & \rho_f & 0 & 0 & m_1 & 0 \\ 0 & 0 & 0 & 0 & \rho_f & 0 & 0 & m_3 \end{pmatrix}. \quad (2.49)$$

We know that \mathbf{E} is a positive-definite matrix because it is the Hessian of the positive-definite quadratic form \mathcal{E} . In fact, because \mathcal{E} has no linear terms in the state variables, we can write it compactly in terms of its Hessian as

$$\mathcal{E} = \frac{1}{2} \mathbf{Q}^T \mathbf{E} \mathbf{Q}. \quad (2.50)$$

For many poroelastic materials, the components of \mathbf{Q} are very badly scaled relative to each other when expressed in common units – for example, waves in the geological materials of Table 4.1 typically have stress components about seven orders of magnitude larger than their velocity components when expressed in SI base units. This makes using the usual vector norms on \mathbf{Q} problematic, but we can fix this issue by using \mathbf{E} to define an energy norm,

$$\|\mathbf{Q}\|_E := \sqrt{\mathbf{Q}^H \mathbf{E} \mathbf{Q}}. \quad (2.51)$$

We define the norm using the Hermitian conjugate-transpose (superscript H) rather than the simple vector transpose in case we later want to take the energy norm of a complex vector. This norm has the advantage of scaling the elements of \mathbf{Q} in a physically relevant fashion, and producing a result that is physically meaningful and has consistent units. Incidentally, this also lets us write the energy density in the even more compact form

$$\mathcal{E} = \frac{1}{2} \|\mathbf{Q}\|_E^2. \quad (2.52)$$

Note that because \mathbf{E} involves the elastic moduli, density, etc. of the medium, it is different for different materials in a heterogeneous domain. Because the energy norm is derived from the energy density at a point, however, energy norms computed in different materials can still be meaningfully compared.

2.7. Hyperbolicity. Equipped with the matrix \mathbf{E} , we can now prove that the left-hand side of (2.42) forms a hyperbolic system.

First, consider the system formed by setting the left-hand side of (2.42) equal to zero,

$$\partial_t \mathbf{Q} + \mathbf{A} \partial_x \mathbf{Q} + \mathbf{B} \partial_z \mathbf{Q} = 0. \quad (2.53)$$

This system is hyperbolic if the linear combination $\check{\mathbf{A}} = n_x \mathbf{A} + n_z \mathbf{B}$ is diagonalizable with real eigenvalues for all real n_x and n_z . Suppose \mathbf{v} is an eigenvector of $\check{\mathbf{A}}$ with eigenvalue λ . Then

$$\check{\mathbf{A}} \mathbf{v} = \lambda \mathbf{v}, \quad (2.54)$$

or, multiplying by \mathbf{E} ,

$$\mathbf{E} \check{\mathbf{A}} \mathbf{v} = \lambda \mathbf{E} \mathbf{v}. \quad (2.55)$$

Since \mathbf{E} is nonsingular, any pair (\mathbf{v}, λ) that satisfies the generalized eigenproblem (2.55) also satisfies the original eigenproblem (2.54).

It is not obvious how this transformation is helpful, but if we examine the component matrices \mathbf{EA} and \mathbf{EB} of $\mathbf{E}\tilde{\mathbf{A}}$, after substantial algebra we can discover they are symmetric:

$$\mathbf{EA} = \begin{pmatrix} 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \quad \mathbf{EB} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \end{pmatrix}. \quad (2.56)$$

Thus $\mathbf{E}\tilde{\mathbf{A}}$ is a symmetric matrix, and (2.55) is a real symmetric generalized eigenproblem with a positive-definite matrix on its right-hand side. As such, it has purely real eigenvalues and a full set of linearly independent eigenvectors. $\tilde{\mathbf{A}}$ is therefore diagonalizable and has pure real eigenvalues, which means that (2.53) is a hyperbolic system.

2.8. Entropy function. We can also show that the energy density is a strictly convex entropy function of the system (2.42), in a sense similar to that of Chen, Levermore, and Liu [16]. Adapting the definition of [16] to the notation used here, a function $\Phi : \mathbb{R}^8 \rightarrow \mathbb{R}$ is a strictly convex entropy function for the system (2.42) if it satisfies the following conditions:

1. $\Phi''(\mathbf{Q})(n_x \mathbf{A} + n_z \mathbf{B})$ is symmetric for all scalars n_x and n_z
2. $(\Phi'(\mathbf{Q}))^T \mathbf{DQ} \leq 0$ for all $\mathbf{Q} \in \mathbb{R}^8$
3. For $\mathbf{Q} \in \mathbb{R}^8$, the following are equivalent:
 - (a) $\mathbf{DQ} = 0$
 - (b) $(\Phi'(\mathbf{Q}))^T \mathbf{DQ} = 0$
4. $\Phi''(\mathbf{Q})$ is positive-definite

Here the primes indicate gradients with respect to \mathbf{Q} , so Φ'' is the Hessian of Φ with respect to \mathbf{Q} . The definition of Chen, Levermore, and Liu includes an additional clause in item 3 related to an operator we call $\mathbf{\Pi}$ (\mathbf{Q} in their notation) that maps from \mathbf{Q} to the conserved quantities of the relaxation part of the system, $\partial_t \mathbf{Q} = \mathbf{DQ}$ — namely, that conditions 3(a) and 3(b) should also be equivalent to $\Phi'(\mathbf{Q})^T = \mathbf{v}^T \mathbf{\Pi}$ for some appropriately-sized vector \mathbf{v} . Rather than take this as part of the definition of a strictly convex entropy function, it is more convenient here to take it as a requirement on $\mathbf{\Pi}$.

Suppose $\Phi(\mathbf{Q}) = \mathcal{E}(\mathbf{Q}) = \frac{1}{2} \mathbf{Q}^T \mathbf{EQ}$. From the preceding sections we already know that conditions 1 and 4 are satisfied, so it only remains to prove conditions 2 and 3. Since $\Phi'(\mathbf{Q}) = \mathbf{EQ}$, condition 2 reduces to $\mathbf{Q}^T \mathbf{EDQ} \leq 0$ for all $\mathbf{Q} \in \mathbb{R}^8$. Using equations (2.46) and (2.49) for \mathbf{D} and \mathbf{E} , and recalling $\Delta_i = \rho m_i - \rho_f^2$, we find that

$$\mathbf{ED} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -\eta/\kappa_1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -\eta/\kappa_3 \end{pmatrix}. \quad (2.57)$$

By inspection, \mathbf{ED} is a symmetric negative-semidefinite matrix, so $\Phi'(\mathbf{Q})^T \mathbf{DQ} \leq 0$ for all $\mathbf{Q} \in \mathbb{R}^8$ and condition 2 is satisfied.

For condition 3, note that 3(a) implies 3(b) since if $\mathbf{DQ} = 0$, necessarily $\Phi'(\mathbf{Q})^T \mathbf{DQ} = 0$. To see that $\Phi'(\mathbf{Q})^T \mathbf{DQ} = \mathbf{Q}^T \mathbf{EDQ} = 0$ implies $\mathbf{DQ} = 0$, note that from (2.57), we have $\mathbf{Q}^T \mathbf{EDQ} = -\frac{\eta}{\kappa_1} q_x^2 - \frac{\eta}{\kappa_3} q_z^2 = 0$ if and only if $q_x = q_z = 0$. Since \mathbf{D} only has nonzero entries in the columns corresponding to q_x and q_z , $\Phi'(\mathbf{Q})^T \mathbf{DQ} = 0$ if and only if $\mathbf{DQ} = 0$. Therefore condition 3 holds, and \mathcal{E} is a strictly convex entropy function for the poroelastic system (2.42).

3. Finite volume solution method.

3.1. Wave propagation. We solve the equations of poroelasticity using a Cartesian grid finite volume approach. This section describes the basics of the finite volume method used here, as well as specifics of how we apply this method to poroelasticity. For a comprehensive discussion of this class of finite volume methods, see LeVeque’s book [36].

The class of finite volume method we use here updates cell averages at every step by solving a Riemann problem between each pair of adjacent grid cells. Thinking of one cell as the “left” cell of the problem, and the other as the “right” cell, the Riemann solution process produces the left-going and right-going fluctuations $\mathcal{A}^-\Delta\mathbf{Q}$ and $\mathcal{A}^+\Delta\mathbf{Q}$ — the changes in the cell variables \mathbf{Q} caused by the left-going and right-going waves — along with a set of waves \mathcal{W}_i with speeds s_i that are used to implement higher-order correction terms. With these correction terms included, the methods used here are second-order accurate. Where solutions are not smooth, wave limiters can be used on the higher-order terms to prevent spurious oscillations. While limiters can reduce the asymptotic order of accuracy of the solution, they often decrease the actual value of the error, depending on the norm being used to measure it and on the grid resolution. They can also improve the qualitative behavior of the solution by suppressing dispersive errors, leading to improved estimates of quantities such as wave arrival times, and keeping total variation from increasing.

For a homogeneous first-order hyperbolic system such as (2.53), the left-going and right-going fluctuations are related to the waves and wave speeds by

$$\mathcal{A}^+\Delta\mathbf{Q} = \sum_{s_i > 0} s_i \mathcal{W}_i, \quad \mathcal{A}^-\Delta\mathbf{Q} = \sum_{s_i < 0} s_i \mathcal{W}_i. \quad (3.1)$$

For a linear problem, such as linear poroelasticity, the waves are simply eigenvectors of the flux Jacobian matrix (for instance, the matrix \mathbf{A} of (2.44) for waves propagating in the x direction in a material with its principal axes aligned with the coordinate axes) associated with the material through which the wave propagates — that is, they have the form $\mathcal{W}_i = \beta_i \mathbf{r}_i$, where \mathbf{r}_i is the eigenvector and β_i is a scalar that gives the strength of the wave. Each wave speed s_i is the corresponding eigenvalue of the flux Jacobian.

A quantity of critical importance in these solution methods is the *CFL number* ν . Informally, the CFL number is the ratio of the distance a wave travels in one timestep to the width of a grid cell; more formally, for a Cartesian grid the global CFL number is

$$\nu = \max_{\text{all cells, all waves}} \max \left(\frac{|s_x|\Delta t}{\Delta x}, \frac{|s_z|\Delta t}{\Delta z} \right). \quad (3.2)$$

Here s_x and s_z are the speeds of waves generated from the Riemann problems in the x and z directions, Δx and Δz are the grid spacings, and Δt is the time step size. The methods used here are stable for $\nu \leq 1$; since ν comes from a maximum over all waves, this means that our stability is limited by the fast P wave.

Because the poroelasticity equations we use here are a linear system, solution of the Riemann problem is straightforward. There is one complication, however. We wish to consider domains composed of multiple materials — in fact, our code has the capability for each grid cell to be made of a different material — so the coefficient matrices are only piecewise constant. We always choose the grid boundaries to coincide with the material boundaries, but we must still solve Riemann problems between domains with different coefficient matrices \mathbf{A} and \mathbf{B} . Suppose we are solving a Riemann problem in the x -direction, with $\mathbf{A} = \mathbf{A}_l$ in the left cell and $\mathbf{A} = \mathbf{A}_r$ in the right cell. We know that in the left cell, the Riemann solution consists of waves with strength β_{li} in the directions of the eigenvectors \mathbf{r}_{li} of \mathbf{A}_l , corresponding to the negative eigenvalues of \mathbf{A}_l ; similarly, in the right cell we will have waves with strength β_{ri} in the directions of eigenvectors \mathbf{r}_{ri} or \mathbf{A}_r , corresponding to the positive eigenvalues of \mathbf{A}_r . There will also be a stationary discontinuity at the cell interface, which will lie in the null space of \mathbf{A}_l and \mathbf{A}_r . (The fact that \mathbf{A} has the same null space for any poroelastic material greatly simplifies matters here, and the corresponding eigenvectors \mathbf{r}_4 and \mathbf{r}_5 will not carry a subscript identifying them with the left or right material.) The total jump in \mathbf{Q} across all the waves and the stationary discontinuity must add up to the difference in \mathbf{Q} between the left and right states, $\Delta\mathbf{Q} = \mathbf{Q}_r - \mathbf{Q}_l$, so we require

$$\Delta\mathbf{Q} = \sum_{i=1}^3 \beta_{li} \mathbf{r}_{li} + \sum_{i=4}^5 \beta_{ri} \mathbf{r}_{ri} + \sum_{i=6}^8 \beta_{ri} \mathbf{r}_{ri} =: \tilde{\mathbf{R}}\boldsymbol{\beta}. \quad (3.3)$$

We can thus compute the wave strengths as $\beta = \tilde{\mathbf{R}}^{-1} \Delta \mathbf{Q}$. In practice, since the strength of the stationary discontinuity is never used directly, we never compute β_4 and β_5 . The same analysis holds for a Riemann problem in the z -direction. This approach corresponds to an open-pore condition between the two poroelastic media, as described by Deresiewicz and Skalak [25] and validated by Gurevich and Schoenberg [32]. Other interface conditions are possible in poroelasticity, such as the closed or partially open pore conditions of Deresiewicz and Skalak, or the loose contact condition of Sharma [44]; we do not model these in this work, but they would be straightforward to incorporate into the Riemann solution process.

Because the eigenstructure of poroelasticity is somewhat complex, we do not compute the eigensystems of \mathbf{A} and \mathbf{B} analytically. Instead, we use LAPACK [1] to compute the eigenvalues and eigenvectors of \mathbf{A} and \mathbf{B} for every poroelastic material present in the model, and the $\tilde{\mathbf{R}}^{-1}$ matrices for each Riemann solve direction and every pair of left and right materials that could occur. For efficiency, we pre-compute these quantities for all materials used (or all possible pairs of materials in the case of $\tilde{\mathbf{R}}^{-1}$) before starting the solution proper, and look them up using a material number stored with each cell during the Riemann solves.

Boundary conditions were implemented using the usual ghost cell approach [36]. We set ghost cell values using either zero-order extrapolation, for boundaries where waves should flow outward and not return, or by setting the ghost cell values equal to the exact solution at the centers of those cells, when we verified our code against known analytic solutions.

3.2. Operator splitting. We include the dissipative part $\mathbf{Q}_t = \mathbf{D}\mathbf{Q}$ of the poroelasticity equations using operator splitting. Since the \mathbf{D} matrix is constant, we can use the exact solution operator $\exp(\mathbf{D}\Delta t)$ to advance the solution by a time increment Δt ; not only is this the most accurate solution available for this part of the system, it is also unconditionally stable and allows the time step to be chosen based solely on stability for the wave propagation part of the system.

The software framework we use offers either Godunov or Strang splitting as a run-time option. Godunov splitting is formally first-order accurate in time and uses a single full-length step of the source term operator per time step, while Strang splitting is second-order and uses two half-steps of the source term. For many practical problems Godunov splitting is a good choice because it displays similar error to Strang — the coefficient of the first-order error term is often small — while being less computationally intensive. However, we primarily use Strang splitting here because it displays substantially greater accuracy for the particular poroelasticity problems we solve, and because our source term is computationally cheap compared to the wave propagation part of the system. For comparison, we also show results for Godunov splitting.

3.3. Stiff regime and subcharacteristic condition. For some cases we consider, the time step is much larger than the characteristic time scales associated with the solution of $\partial_t \mathbf{Q} = \mathbf{D}\mathbf{Q}$. These cases fall outside the regime where asymptotic leading-order error estimates are relevant, and for them the source term is stiff, a known source of difficulty in operator splitting approaches for hyperbolic equations [18, 37]. Based on a conjecture of Pember [41], we expect to avoid spurious solutions for this stiff relaxation system if the poroelasticity equations (2.42) satisfy a *subcharacteristic condition*, where the wave speeds for the reduced equations obtained by restricting the full system to the equilibrium manifold of the dissipation term (i.e. zero fluid velocity relative to the solid matrix) interlace with the wave speeds for the full system. The appropriate generalization of Pember’s conjecture to systems of more than two equations is not obvious, but based on the principle that information should propagate more slowly (certainly no more quickly!) in the reduced system than in the full system, we expect to avoid spurious solutions if for all possible wave propagation directions the speeds λ_1 and λ_2 of the reduced system are strictly less than the speed of the fastest wave of the full system,

$$\lambda_1 < c_{pf}, \quad \lambda_2 < c_{pf}. \quad (3.4)$$

Here c_{pf} is the speed of a fast P wave. We ignore the negative eigenvalues for both the full and reduced systems, because they are simply the negatives of the wave speeds and will automatically satisfy a similar inequality. We also ignore the zero eigenvalues of the full system, since they correspond to eigencomponents of the solution that are left unchanged in the wave propagation part of the solution process, and are evolved according to the exact solution operator $\exp(\mathbf{D}\Delta t)$ in the dissipation part.

To construct the appropriate reduced system, we follow the derivation of Chen, Levermore, and Liu [16]. First we examine the dissipation part of the system in isolation,

$$\partial_t \mathbf{Q} = \mathbf{DQ}. \quad (3.5)$$

System (3.5) has six conserved quantities $\mathbf{u} = \mathbf{\Pi Q}$, which are related to the state variables \mathbf{Q} by the matrix

$$\mathbf{\Pi} := \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & \rho_f/\rho & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & \rho_f/\rho \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \end{bmatrix}. \quad (3.6)$$

The fact that \mathbf{u} is a vector of conserved quantities of (3.5) follows immediately from the fact that $\mathbf{\Pi D} = 0$, so that $\partial_t \mathbf{u} = \mathbf{\Pi} \partial_t \mathbf{Q} = \mathbf{\Pi DQ} = 0$. Given the conserved quantities \mathbf{u} , the unique equilibrium \mathbf{Q}_{eq} of (3.5) that satisfies $\mathbf{DQ}_{\text{eq}} = 0$ and $\mathbf{\Pi Q}_{\text{eq}} = \mathbf{u}$ is $\mathbf{Q}_{\text{eq}} = \mathbf{Gu}$, where

$$\mathbf{G} := \begin{bmatrix} \mathbf{I}_{6 \times 6} \\ \mathbf{0}_{2 \times 6} \end{bmatrix}. \quad (3.7)$$

Notice that $\mathbf{\Pi}$ and \mathbf{G} satisfy the relation $\mathbf{\Pi G} = \mathbf{I}_{6 \times 6}$. The reduced system is found by multiplying the full poroelastic system from the left by $\mathbf{\Pi}$ to eliminate the dissipation term, and requiring the state vector \mathbf{Q} to lie on the equilibrium manifold, $\mathbf{Q} = \mathbf{Q}_{\text{eq}} = \mathbf{Gu}$, resulting in

$$\partial_t \mathbf{u}_t + \mathbf{\Pi A G} \partial_x \mathbf{u} + \mathbf{\Pi B G} \partial_z \mathbf{u} = 0. \quad (3.8)$$

Now that we have the matrix $\mathbf{\Pi}$, we can show that the additional condition of Chen, Levermore, and Liu mentioned in Section 2.8 also holds — that the statements $\mathbf{DQ} = 0$ and $\Phi'(\mathbf{Q})^T \mathbf{DQ} = \mathbf{Q}^T \mathbf{E DQ} = 0$ are equivalent to $\Phi'(\mathbf{Q})^T = \mathbf{Q}^T \mathbf{E} = \mathbf{v}^T \mathbf{\Pi}$ for some $\mathbf{v} \in \mathbb{R}^6$. First, if $\mathbf{Q}^T \mathbf{E} = \mathbf{v}^T \mathbf{\Pi}$, we immediately have $\mathbf{Q}^T \mathbf{E DQ} = \mathbf{v}^T \mathbf{\Pi DQ} = 0$ since $\mathbf{\Pi D} = 0$. We also noted in section 2.8 that if $\mathbf{DQ} = 0$, then $q_x = q_z = 0$. Thus the fourth, fifth, seventh, and eighth components of $\mathbf{Q}^T \mathbf{E}$ are $(\mathbf{Q}^T \mathbf{E})_4 = \rho v_x$, $(\mathbf{Q}^T \mathbf{E})_5 = \rho v_z$, $(\mathbf{Q}^T \mathbf{E})_7 = \rho_f v_x$, and $(\mathbf{Q}^T \mathbf{E})_8 = \rho_f v_z$. We therefore have $\mathbf{Q}^T \mathbf{E} = \mathbf{v}^T \mathbf{\Pi}$ with \mathbf{v} given by

$$\mathbf{v} = \begin{pmatrix} (\mathbf{Q}^T \mathbf{E})_1 \\ (\mathbf{Q}^T \mathbf{E})_2 \\ (\mathbf{Q}^T \mathbf{E})_3 \\ \rho v_x \\ \rho v_z \\ (\mathbf{Q}^T \mathbf{E})_6 \end{pmatrix}. \quad (3.9)$$

Given that a strictly convex entropy function $\Phi(\mathbf{Q}) = \mathcal{E}(\mathbf{Q})$ exists and the above condition is satisfied, Theorem 2.1 and the subsequent remark of Chen, Levermore, and Liu [16] imply that the reduced system (3.8) is hyperbolic, and satisfies a nonstrict subcharacteristic condition, which we can render here in the context of poroelasticity as

$$c_{ps} \leq \lambda_2 \leq c_{pf}, \quad 0 \leq \lambda_1 \leq c_s. \quad (3.10)$$

While this does not imply the strict subcharacteristic condition (3.4), it is nearly as useful. This characterization of the eigenvalues comes from the fact that eigenvalues of the symmetric generalized eigenproblem (2.55) satisfy the Rayleigh quotient minimax principle

$$\lambda_{fk}^\downarrow = \max_{\mathcal{S}_k} \min_{\mathbf{v} \in \mathcal{S}_k} \frac{\mathbf{v}^T \mathbf{E} \check{\mathbf{A}} \mathbf{v}}{\mathbf{v}^T \mathbf{E} \mathbf{v}}, \quad \lambda_{fk}^\uparrow = \min_{\mathcal{S}_k} \max_{\mathbf{v} \in \mathcal{S}_k} \frac{\mathbf{v}^T \mathbf{E} \check{\mathbf{A}} \mathbf{v}}{\mathbf{v}^T \mathbf{E} \mathbf{v}}, \quad (3.11)$$

where \mathcal{S}_k is any k -dimensional subspace of \mathbb{R}^8 , λ_{fk}^\downarrow is the k 'th eigenvalue of the full system counting down from the largest, and λ_{fk}^\uparrow is the k 'th eigenvalue counting up from the smallest. Chen, Levermore, and Liu

prove that the eigenvalues of the reduced system satisfy a similar minimax principle over a restricted set of subspaces,

$$\lambda_{ek}^\downarrow = \max_{\mathcal{S}_k \subseteq \mathcal{R}(\mathbf{G})} \min_{\mathbf{v} \in \mathcal{S}_k} \frac{\mathbf{v}^T \mathbf{E} \check{\mathbf{A}} \mathbf{v}}{\mathbf{v}^T \mathbf{E} \mathbf{v}}, \quad \lambda_{ek}^\uparrow = \min_{\mathcal{S}_k \subseteq \mathcal{R}(\mathbf{G})} \max_{\mathbf{v} \in \mathcal{S}_k} \frac{\mathbf{v}^T \mathbf{E} \check{\mathbf{A}} \mathbf{v}}{\mathbf{v}^T \mathbf{E} \mathbf{v}}, \quad (3.12)$$

where the subscript e indicates eigenvalues of the reduced (“equilibrium”) system and $\mathcal{R}(\mathbf{G})$ is the range space of \mathbf{G} .

Equality can be realized in the nonstrict subcharacteristic condition (3.10) — for example, if the fluid density ρ_f for the orthotropic sandstone material whose parameters are given in the first column of Table 4.1 is reduced from 1040 kg/m³ to 208.9 kg/m³, a fast P wave traveling along the material principal 1-axis shows no fluid relative motion. This means its eigenvector lies in $\mathcal{R}(\mathbf{G})$, so by (3.12) the reduced system shares the same wave speed. Fortunately, this turns out to be innocuous from the standpoint of the true solution — if the eigenvector \mathbf{v} associated with some wave is in $\mathcal{R}(\mathbf{G})$, then it is in $\mathcal{N}(\mathbf{D})$; since it is an eigenvector of both parts of the system, the corresponding eigencomponent of the solution can be decoupled from the rest of the system, and its solution is independent of them. Furthermore, the PDE describing the evolution of this component of the solution is a purely hyperbolic one, with no source term — there is no momentum transfer due to viscous drag between the solid and fluid for this wave mode because there is no relative motion between them. A similar result carries over to the numerical solution obtained by operator splitting: if \mathbf{v} is in $\mathcal{N}(\mathbf{D})$, then $\exp(\mathbf{D}\Delta t)\mathbf{v} = \mathbf{v}$, and the corresponding component of \mathbf{Q} passes through the solution operator for the dissipation term unchanged.

As an aside, the reduced system (3.8) has a familiar form. If we multiply out the coefficient matrices, we get

$$\mathbf{\Pi A G} = \begin{pmatrix} 0 & 0 & 0 & -c_{11}^u & 0 & 0 \\ 0 & 0 & 0 & -c_{13}^u & 0 & 0 \\ 0 & 0 & 0 & 0 & -c_{55}^u & 0 \\ -\frac{1}{\rho} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -\frac{1}{\rho} & 0 & 0 & 0 \\ 0 & 0 & 0 & M\alpha_1 & 0 & 0 \end{pmatrix}, \quad \mathbf{\Pi B G} = \begin{pmatrix} 0 & 0 & 0 & 0 & -c_{13}^u & 0 \\ 0 & 0 & 0 & 0 & -c_{33}^u & 0 \\ 0 & 0 & 0 & -c_{55}^u & 0 & 0 \\ 0 & 0 & -\frac{1}{\rho} & 0 & 0 & 0 \\ 0 & -\frac{1}{\rho} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & M\alpha_3 & 0 \end{pmatrix}. \quad (3.13)$$

The upper-left 5×5 portions of these matrices are just the coefficient matrices for orthotropic plane-strain elasticity, with the fluid pressure coming along as an additional variable determined entirely by the elastic field variables. Because of this we will identify the faster wave of the reduced system as the “reduced P wave” and the slower one as the “reduced S wave.” Note that for this work, the reduced system is only of theoretical importance — for the actual numerical code, we discretize the full system (2.42).

3.4. Numerical software. We implemented the numerical solution techniques described here using the CLAWPACK finite volume method package, version 4.6 [45]. CLAWPACK implements the parts of a high-resolution finite volume code that are common across all problems, leaving the user to write only problem-specific code such as Riemann solvers. Operator splitting is supported for source terms, such as the dissipative term here, by means of a user-supplied subroutine that advances the system by a specified time step under the action of the source term. Both Godunov and Strang splitting are available. Block-structured Berger-Colella-Oliger adaptive mesh refinement (AMR) is available from the AMRCLAW package [3]; AMRCLAW can also run in parallel on shared-memory systems using OpenMP. Besides Cartesian grids, CLAWPACK and AMRCLAW also support logically rectangular mapped grids — while we do not use mapped grids here, we plan to present results with them in a subsequent publication.

4. Results. We present results here for four classes of problems. First, we demonstrate convergence of our numerical solution to known analytic plane wave solutions for an orthotropic medium for the wave propagation part of the system alone. We then include the viscous dissipation term, and examine the effect of operator splitting on accuracy, again comparing against known analytic plane wave solutions. Next, we show results for simple point sources in uniform orthotropic media, solving test problems previously addressed by de la Puente et al. [23] and Carcione [12] in order to further verify our code. Finally, we solve a larger-scale problem with a domain composed of two isotropic materials, involving wave reflection, refraction, and

interconversion at the material boundary, as well as demonstrating the use of adaptive mesh refinement to reduce the time required for solution.

4.1. Analytic plane wave solution. Before we test our code's convergence against analytic plane wave solutions, however, we will first outline the procedure used to obtain these analytic solutions. We start by assuming the velocity and stress fields have a plane wave form,

$$\mathbf{V} := (v_x \quad v_z \quad q_x \quad q_z)^T = \mathbf{V}_0 \exp(i(k_x x + k_z z - \omega t)) \quad (4.1)$$

$$\mathbf{T} := (\tau_{xx} \quad \tau_{zz} \quad \tau_{xz} \quad -p)^T = \mathbf{T}_0 \exp(i(k_x x + k_z z - \omega t)). \quad (4.2)$$

Here \mathbf{V}_0 and \mathbf{T}_0 are constant vectors, and ω is the prescribed angular frequency of the wave. The wavenumbers k_x and k_z are yet to be determined, but we also prescribe $k_x = kl_x$ and $k_z = kl_z$, where l_x and l_z are the (real-valued) direction cosines of the wavevector, with $l_x^2 + l_z^2 = 1$.

With these assumptions on the solution, stress-strain equations (2.30) through (2.33) imply

$$-\omega \mathbf{T}_0 = k \mathbf{F} \mathbf{V}_0, \quad (4.3)$$

where the matrix \mathbf{F} is

$$\mathbf{F} = \begin{pmatrix} l_x c_{11}^u & l_z c_{13}^u & \alpha_1 M l_x & \alpha_1 M l_z \\ l_x c_{13}^u & l_z c_{33}^u & \alpha_3 M l_x & \alpha_3 M l_z \\ l_z c_{55}^u & l_x c_{55}^u & 0 & 0 \\ \alpha_1 M l_x & \alpha_3 M l_z & M l_x & M l_z \end{pmatrix}. \quad (4.4)$$

Equations of motion (2.34) through (2.37) also imply

$$k \mathbf{L} \mathbf{T}_0 = -\omega \mathbf{\Gamma} \mathbf{V}_0, \quad (4.5)$$

where the matrices \mathbf{L} and $\mathbf{\Gamma}$ are

$$\mathbf{L} = \begin{pmatrix} l_x & 0 & l_z & 0 \\ 0 & l_z & l_x & 0 \\ 0 & 0 & 0 & l_x \\ 0 & 0 & 0 & l_z \end{pmatrix}, \quad \mathbf{\Gamma} = \begin{pmatrix} \rho & 0 & \rho_f & 0 \\ 0 & \rho & 0 & \rho_f \\ \rho_f & 0 & iY_1(-\omega)/\omega & 0 \\ 0 & \rho_f & 0 & iY_3(-\omega)/\omega \end{pmatrix}, \quad (4.6)$$

and $Y_j(\omega) := i\omega m_j + \eta/\kappa_j$ for $j = 1, 3$.

Combining equations (4.3) and (4.5), we can obtain an eigenproblem for \mathbf{V}_0 and $(\frac{\omega}{k})^2$:

$$\mathbf{\Gamma}^{-1} \mathbf{L} \mathbf{F} \mathbf{V}_0 = \left(\frac{\omega}{k}\right)^2 \mathbf{V}_0. \quad (4.7)$$

To obtain a plane wave solution, we solve this eigenproblem, choose the eigenvalue $(\frac{\omega}{k})^2$ and eigenvector \mathbf{V}_0 corresponding to the wave family of interest, then back out the appropriate \mathbf{T}_0 from (4.3). Note that the eigenvalues and eigenvectors will be complex-valued if dissipation is present.

4.2. Plane wave convergence results — inviscid. We conducted our plane wave convergence tests on a uniform domain composed of orthotropic layered sandstone, whose properties are given in Table 4.1. The wave-propagation part of our code was first tested alone, without viscous dissipation included. These tests were conducted using plane waves with a fixed angular frequency of 10^4 rad/s, in a square domain 8 m on a side. This distance is about two wavelengths of the fast P wave at this frequency. The total simulation time was $2\pi \times 10^{-4}$ s — one period of the wave. Because there is no intrinsic time scale associated with Biot's equations when viscous dissipation is omitted, the frequency of the plane wave is not directly relevant to the accuracy of these tests; instead, only the number of grid cells per wavelength and the ratio of the simulation time to the wave period are relevant. Boundary conditions were set by filling the ghost cells with the value of the true plane wave solution at the cell centers, and the initial condition on the grid was set the same way. This ensured that the accuracy of the numerical solution was governed only by the correctness of the

Table 4.1: Properties of the poroelastic media used in test cases, taken from de la Puente et al. [23]. Wave speeds are correct in the high-frequency limit. c_{pf} is the fast P wave speed, c_s is the S wave speed, c_{ps} is the slow P wave speed, and τ_d is the time constant for dissipation. Subscript numbers indicated principal directions.

		Sandstone (orthotropic)	Glass/epoxy (orthotropic)	Sandstone (isotropic)	Shale (isotropic)
Base properties					
K_s	(GPa)	80	40	40	7.6
ρ_s	(kg/m ³)	2500	1815	2500	2210
c_{11}	(GPa)	71.8	39.4	36	11.9
c_{12}	(GPa)	3.2	1.2	12	3.96
c_{13}	(GPa)	1.2	1.2	12	3.96
c_{33}	(GPa)	53.4	13.1	36	11.9
c_{55}	(GPa)	26.1	3.0	12	3.96
ϕ		0.2	0.2	0.2	0.16
κ_1	(10 ⁻¹⁵ m ²)	600	600	600	100
κ_3	(10 ⁻¹⁵ m ²)	100	100	600	100
T_1		2	2	2	2
T_3		3.6	3.6	2	2
K_f	(GPa)	2.5	2.5	2.5	2.5
ρ_f	(kg/m ³)	1040	1040	1040	1040
η	(10 ⁻³ kg/m.s)	1	1	0	0
Derived quantites					
c_{pf1}	(m/s)	6000	5240	4250	2480
c_{pf3}	(m/s)	5260	3580	4250	2480
c_{s1}	(m/s)	3480	1370	2390	1430
c_{s3}	(m/s)	3520	1390	2390	1430
c_{ps1}	(m/s)	1030	975	1020	1130
c_{ps3}	(m/s)	746	604	1020	1130
τ_{d1}	(μ s)	5.95	5.85	—	—
τ_{d3}	(μ s)	1.82	1.81	—	—

wave-propagation algorithm used within the problem domain, not by the implementation of the boundary conditions.

Since we are working with an orthotropic material, rather than an isotropic one, the speed and associated eigenvector for a plane wave depend on its propagation direction — the solutions for plane waves propagating in different directions are not simply rotated versions of each other, and in order to be confident in the correctness of our code it was necessary to test it with plane waves propagating at a variety of angles θ_{wave} relative to the global x axis. Since we anticipate solving problems where the principal material axes do not coincide with the global coordinate axes, we also tested our code with a variety of angles θ_{mat} between the material 1 axis and the x axis. Figure 4.1 shows a sample plane wave solution, with the relevant axes and angles identified. For the studies presented here, we used θ_{wave} values from 0° to 345° counterclockwise, and θ_{mat} values from 0° to 165° counterclockwise, both in steps of 15°. (When expressed in global x - z coordinates, the system matrices \mathbf{A} and \mathbf{B} simply flip sign after a rotation of $\theta_{\text{mat}} = 180^\circ$, while \mathbf{D} is unchanged, so it was not necessary to use θ_{mat} values of 180° or over.) This gives 24 different θ_{wave} values and 12 different θ_{mat} values, for a total of 288 combinations of these angles. For each $(\theta_{\text{wave}}, \theta_{\text{mat}})$ pair, we examined plane waves in each of the three families, on 100×100 , 200×200 , 400×400 , and 800×800 cell grids, for a total of 3456 different test cases. For all cases, we chose the time step so that the CFL number was 0.9. Because the solution is smooth, we used no wave limiting for this convergence study.

For each test case, we measured the error by taking the energy norm on each cell of the difference between




Fig. 4.1: Sample plane wave initial condition, showing the relation between the global $x-z$ axes, the material 1 – 3 axes, and the wave propagation direction. For this plot $\theta_{\text{mat}} = 15^\circ$ and $\theta_{\text{wave}} = -30^\circ$. The shading shows the z -direction solid matrix velocity. This plot depicts a fast P wave, with an angular frequency of 10^4 rad/s, in the same orthotropic sandstone medium used for the inviscid convergence tests. The grid dimensions are 800×800 cells.

the numerically obtained cell value and the true solution at the cell center. We then applied the grid 1-, 2-, and max-norms to the energy norm error field to obtain an aggregate error norm. (The grid 1-norm used here is just the linear algebraic 1-norm divided by the number of grid cells; similarly, the grid 2-norm is the linear algebraic 2-norm divided by the square root of the number of grid cells.) For each combination of θ_{wave} , θ_{mat} , and wave family, we performed a linear least-squares fit of $\log(\text{error})$ versus $\log(m)$, where m is the number of grid cells along each axis. The slope of this fit was considered to be the convergence rate of the code for this set of cases.

Table 4.2 summarizes the results of this convergence study. The convergence rates listed are the maximum, minimum, and mean over all combinations of θ_{wave} and θ_{mat} ; the worst R^2 value for the least-squares fit of $\log(\text{error})$ versus $\log(m)$ is also reported. The last two columns of Table 4.2 give the lowest and highest error for each wave in each norm on the finest grid, observed over all combinations of θ_{wave} and θ_{mat} , normalized by the energy norm of the plane wave eigenvector \mathbf{Q}_0 formed from the vectors \mathbf{V}_0 and \mathbf{T}_0 of the preceding section in order to allow a fair comparison between different cases. We see second-order convergence in all three norms for the fast P and S waves. Results are impaired for the slow P-wave because its slow propagation speed causes it to be underresolved on the coarser grids at the frequency used, but we still observe second-order convergence in the 1-norm and 2-norm. While the error varies somewhat depending on the wave propagation and principal material directions, it does so by no more than a factor of 3-4, indicating that there are no severe grid alignment effects.

We do not include convergence results with wave limiting here, but informal exploration suggests that for the cases above, limiting reduces the order of accuracy to about 1.9 in the 1-norm, 1.8 in the 2-norm, and 1.6 in the max-norm. This is because wave limiting tends to clip extrema in order to avoid introducing spurious oscillations. Despite the reduced order of accuracy, however, using a limiter can improve actual error in many cases, often in the 1-norm but even in the max-norm if a wave is poorly resolved or heavily affected by dispersive errors. Figure 4.2 shows an example of the effect of limiting, using the Monotonized Centered (MC) limiter, on the normalized energy max-norm and 1-norm errors in each of the three wave families for the inviscid test cases with $\theta_{\text{mat}} = \theta_{\text{wave}} = 0$. For each wave, the max-norm error decreases more slowly with increasing grid size when the limiter is present. The fast P wave is well resolved even on the coarsest grid, and always shows lower max-norm error without limiting. The S wave is somewhat less well-resolved, and the difference between the two curves is smaller, with the max-norm errors both with and without limiters roughly equal on the coarsest grid. Finally, the slow P wave starts out poorly resolved, and

Table 4.2: Summary of convergence results for inviscid test cases

Wave family	Error norm	Convergence rate			Worst R^2 value	Error on 800×800 grid	
		Best	Worst	Mean		Best	Worst
Fast P	1-norm	2.03	2.01	2.01	0.99994	2.53×10^{-5}	6.90×10^{-5}
	2-norm	2.02	2.01	2.01	0.99995	3.04×10^{-5}	8.23×10^{-5}
	Max-norm	2.02	1.96	2.00	0.99943	5.61×10^{-5}	1.76×10^{-4}
S	1-norm	2.01	2.00	2.01	1.00000	1.31×10^{-4}	3.23×10^{-4}
	2-norm	2.01	2.00	2.00	1.00000	1.49×10^{-4}	3.74×10^{-4}
	Max-norm	2.00	1.96	1.99	0.99977	2.80×10^{-4}	7.98×10^{-4}
Slow P	1-norm	1.99	1.92	1.97	0.99922	3.05×10^{-3}	1.13×10^{-2}
	2-norm	1.99	1.94	1.97	0.99953	3.44×10^{-3}	1.27×10^{-2}
	Max-norm	1.93	1.67	1.80	0.99146	8.81×10^{-3}	3.16×10^{-2}




Fig. 4.2: Comparison of error both with and without a limiter for the inviscid plane wave test cases with $\theta_{\text{mat}} = \theta_{\text{wave}} = 0$. The circles indicate normalized energy max-norm error, while the stars indicate normalized energy 1-norm error. The limiter is always beneficial in the 1-norm for these cases, and is also helpful in the max-norm for poorly-resolved waves.

using a limiter produces lower max-norm error on all but the finest grid. The 1-norm error is lower with the limiter included for all cases. For further discussion of the benefits and drawbacks of using limiters, see LeVeque [36]. There have also been efforts to produce limiters that are compatible with higher-order methods; see for example Čada and Torrilhon [11], Liu and Tadmor [38], or the recent review by Kemm [35].

4.3. Plane wave convergence results — viscous. With viscosity included, numerical solution of the equations of poroelasticity becomes substantially more challenging. The chief difficulty is that the dissipation term has its own associated time scales, independent of the computational grid. Since an appropriate time




Fig. 4.3: Wave speeds (meters per second) for full and reduced systems for both materials used with viscosity included, as a function of propagation angle. Zero propagation angle corresponds to the material principal 1-axis. The strict subcharacteristic condition (3.4) is satisfied.

step for the wave propagation part of the system is proportional to the grid size — preferably with a CFL number near 1 — for large enough grid cell sizes the dissipation term is stiff relative to the wave propagation term. While we maintain stability by solving the dissipation term exactly, because the wave propagation and dissipation parts of the system do not commute, we can still heuristically expect problems in our operator splitting scheme if the time step is much longer than the characteristic time scale for dissipation.

Revisiting the subcharacteristic condition of section 3.3, Figure 4.3 shows the wave speeds for the full and reduced systems as a function of propagation direction relative to the principal axes. The strict subcharacteristic condition (3.4) is satisfied for all the materials we examine, although the reduced P wave speed nearly reaches the fast P wave speed for the glass/epoxy material. In fact, an even stricter condition is satisfied: the wave speeds interleave, with exactly one wave of the reduced system between each consecutive pair of waves of the full system. Based on the discussion of section 3.3, this suggests that we will not see spurious solutions or incorrect wave speeds from our numerical solution.

To better explore the behavior of our numerical method in the presence of dissipation, we ran a series of numerical tests in the same sandstone medium as the inviscid test cases, against plane wave solutions at frequencies ranging from 10 Hz to 20 kHz. (The maximum frequency for low-frequency Biot theory to be valid in this medium is roughly 25 kHz.) For all cases, the material 1-axis was aligned with the global x -axis, and waves were set to propagate in the positive x direction. For the fast P and S waves, we chose the domain size to be two damped wavelengths of the wave in question; since the slow P wave has a characteristic decay length (the distance over which the wave amplitude decreases by a factor of e) that is typically a fifth or less of its wavelength, we chose the domain size for the slow P wave cases to be twice the characteristic decay length instead. We chose the domain sizes in this way so that the number of grid cells per wavelength, or per decay length, would be constant across all frequencies; this keeps the discretization error contributed from the wave propagation part of the system roughly constant for each grid size across all frequencies, helping to isolate the error caused by operator splitting. The total simulation time for the fast P and S wave cases was 1.25 cycles of the wave (we chose a non-integer number of cycles to avoid any possible spoofing where an unchanged solution might appear correct), while for the slow P wave cases it was 1.25 times the time for a fast P wave traveling in the x direction to cross the domain. For each combination of wave family and frequency, we computed solutions on grids of size 100×100 , 200×200 , 400×400 , and 800×800 , using both Godunov and Strang splitting.

Figure 4.4 shows the results of these tests in the same normalized energy max-norm used for the inviscid cases. There is a pronounced qualitative difference in convergence behavior depending on frequency. At low




Fig. 4.4: Normalized energy max-norm error for operator splitting tests. For each splitting method (Godunov or Strang) within each subplot, the curves fall in order of increasing grid fineness, with each curve from a grid twice as fine as the curve above. The top curve is on a 100×100 grid, and the bottom is on an 800×800 . The circles indicate where the time step was less than the characteristic time scale τ_d for dissipation in the x direction; the stars indicate where the time step was greater.

frequencies, corresponding to large grid cell sizes and long time steps, both splitting methods show first-order convergence. Starting at a step length of roughly 5-10 times the characteristic time τ_d for dissipation in the x direction (the time over which the x velocity of the fluid relative to the matrix decreases by a factor of e), the two methods begin behaving differently, with the Godunov splitting error increasing abruptly while the Strang splitting error sweeps smoothly down to second-order convergence. This effect is most visible for the fast P wave, but can also be seen slightly in the 400×400 grid and strongly in the 800×800 grid for the S wave. Because of the choice of domain size, the slow P test cases always had a time step below the characteristic dissipation time; they display consistent first-order convergence with Godunov splitting and second-order with Strang. The qualitative shift in behavior with time step length can be understood by noting that for a time step much longer than the characteristic time, the solution operator $\exp(\mathbf{D}\Delta t)$ of the dissipative part of the system is essentially a projection operator that sets the fluid relative velocity to zero and transfers all of the fluid relative momentum into the bulk motion of the medium. The effect of this projection operator is essentially the same whether it is applied once per time step after solution of the wave propagation part of the system (Godunov splitting), or twice, both before and after wave propagation (Strang splitting).

With these results available to inform our choices, we conducted a set of convergence studies similar to the inviscid cases in the previous subsection. Because of the qualitative difference in convergence behavior for different time step regimes, we performed convergence studies both at a point in the high-frequency convergence regime of Figure 4.4 (10 kHz), and at a low-frequency point (10 Hz). Since the slow P wave decays extremely rapidly in the presence of viscosity — typically by a factor of 10 to 100 or more per wavelength in the valid frequency range for Biot theory — we chose a different domain size for the viscous

Table 4.3: Summary of convergence results for viscous, high-frequency test cases

Wave family	Error norm	Convergence rate			Worst R^2 value	Error on 800×800 grid	
		Best	Worst	Mean		Best	Worst
Fast P	1-norm	2.03	2.01	2.01	0.99996	2.61×10^{-5}	7.31×10^{-5}
	2-norm	2.03	2.01	2.01	0.99996	3.16×10^{-5}	8.79×10^{-5}
	Max-norm	2.05	2.00	2.02	0.99887	6.44×10^{-5}	1.77×10^{-4}
S	1-norm	2.01	2.01	2.01	1.00000	1.38×10^{-4}	3.39×10^{-4}
	2-norm	2.01	2.00	2.01	1.00000	1.58×10^{-4}	3.98×10^{-4}
	Max-norm	2.03	1.99	2.00	0.99978	2.99×10^{-4}	7.59×10^{-4}
Slow P	1-norm	2.01	2.00	2.01	1.00000	2.12×10^{-6}	1.74×10^{-4}
	2-norm	2.02	2.00	2.01	1.00000	2.49×10^{-6}	3.26×10^{-4}
	Max-norm	2.02	1.96	2.00	0.99988	6.53×10^{-6}	2.25×10^{-3}

slow P test cases. For the high-frequency cases, we chose square domains with side length 1.2 meters for the fast P and S waves — roughly two wavelengths of the fast P wave at 10 kHz — and 5 centimeters for the slow P wave — roughly 2-5 times the characteristic decay length for this wave, depending on propagation direction in the orthotropic medium. For the low-frequency cases, the domain size was 1200 meters for the fast P and S waves, and 1 meter for the slow P wave — again, 2-4 times the characteristic decay length, which is far shorter than the wavelength. This huge disparity in domain sizes is somewhat troublesome, but it is not clear whether simulation results for a slow P wave on a domain of the size used for the fast P and S waves would be meaningful, since the solution decays over such a short distance. For practical problems, this would be an excellent opportunity for adaptive mesh refinement, to generate fine grids where and when slow P waves appear, then coarsen the grid again after they dissipate.

We ran all of the viscous test cases to essentially the same final times as for the frequency sweep of Figure 4.4. For the 10 kHz runs, this was $125 \mu\text{s}$ for the fast P and S waves (1.25 periods of the wave), and $10.4 \mu\text{s}$ for the slow P wave (1.25 times the time for a fast P wave to cross the domain), while for the 10 Hz runs it was 0.125 s for the fast P and S waves, and $208 \mu\text{s}$ for the slow P wave. Strang splitting was used for all cases. All other aspects of the solution, including the sets of wave propagation and material principal directions θ_{wave} and θ_{mat} as well as the method of setting the boundary conditions, were the same as for the inviscid test cases, and we measured the solution error using the same set of norms.

Table 4.3 summarizes convergence for the high-frequency viscous cases. We again see consistent second-order convergence in all norms for the fast P and S waves at high frequency, and a similar amount of dependence of error on wave propagation direction. For the slow P wave at high frequency, however, results are substantially different from the inviscid cases. We see second-order convergence, since the solution is well-resolved on the grids used here, but there is now a factor of several hundred difference between the maximum and minimum error at a single grid size. Close examination of the error for individual cases shows that it is primarily a function of the offset $\theta_{\text{wave}} - \theta_{\text{mat}}$; for a fixed value of $\theta_{\text{wave}} - \theta_{\text{mat}}$ and a fixed grid size, the error is similar across all values of θ_{mat} . This indicates that the large variation in error is a effect of the alignment of the wavefront relative to the principal material axes, rather than a grid alignment effect. The likely cause is the substantial difference in the characteristic decay times between the 1 and 3 axes of the material — the decay time in the 1 direction is 5.95 microseconds, while in the 3 direction it is 1.82 microseconds. This large variation in decay time causes a large variation in the operator splitting error. In addition, the characteristic decay length is substantially shorter in the 3 direction, causing the solution magnitude to be larger at the “upstream” (opposite the propagation direction) edge of the domain relative to the value at the center against which the error is normalized; the larger solution magnitude naturally results in a larger error.

Table 4.4 shows the results for the low-frequency viscous test cases. The fast P and S waves again show only a weak dependence of error on wave propagation and principal material direction, but their convergence rates are substantially degraded, just as in the low-frequency range of Figure 4.4. Convergence of the fast P wave is roughly first-order in the worst case, due to the long time step relative to the characteristic decay

Table 4.4: Summary of convergence results for viscous, low-frequency test cases

Wave family	Error norm	Convergence rate			Worst R^2 value	Error on 800×800 grid	
		Best	Worst	Mean		Best	Worst
Fast P	1-norm	1.57	1.10	1.26	0.99239	1.95×10^{-4}	4.24×10^{-4}
	2-norm	1.60	1.10	1.28	0.99206	2.15×10^{-4}	4.77×10^{-4}
	Max-norm	1.61	1.10	1.32	0.99014	3.90×10^{-4}	1.04×10^{-3}
S	1-norm	1.85	1.37	1.58	0.99241	4.37×10^{-4}	7.30×10^{-4}
	2-norm	1.85	1.37	1.59	0.99274	5.01×10^{-4}	8.30×10^{-4}
	Max-norm	1.83	1.36	1.58	0.98917	9.63×10^{-4}	1.92×10^{-3}
Slow P	1-norm	2.00	1.94	1.98	0.99985	1.04×10^{-6}	4.88×10^{-5}
	2-norm	2.03	1.93	1.99	0.99982	1.08×10^{-6}	7.40×10^{-5}
	Max-norm	2.09	1.84	1.98	0.99795	2.40×10^{-6}	4.59×10^{-4}

times. Surprisingly, results for the S wave are only slightly better in the worst case, likely due to the much shorter characteristic decay time in the 3 direction. If the grid were further refined, we would presumably reach a second-order convergence regime as the time step approached the characteristic decay time, but in order to have a time step similar to the shortest decay time at a CFL number of 0.9 we would need a grid cell size of roughly 13 cm — resulting in a 90000×90000 cell grid on the 1200 m square domain if this size is uniform! The low-frequency slow P wave cases, by contrast, show the same strong dependence of error on the alignment of the wave direction with the principal material direction as the high-frequency cases, but because the domain size was much smaller and the time step much shorter, we observe consistent second-order convergence.

As a final comment on convergence, we note that while the rate of reduction of error with decreasing mesh size is poorer for the low-frequency fast P and S wave cases, this may not be a problem in practice unless very high accuracy is desired. For both frequency ranges investigated, the relative error in all three norms on the 800×800 grid never exceeds 1.04×10^{-3} for the fast P wave, or 1.92×10^{-3} for the S wave — despite the numerical difficulties encountered, the waves are still well-resolved.

4.4. Single-material point source results. With the accuracy of our method characterized for simple plane wave solutions, we are able to move on to more interesting problems. First, we compare against the results of de la Puente et al. [23] and Carcione [12] for a point source in a uniform orthotropic medium. We used the orthotropic sandstone and glass/epoxy media described in Table 4.1; in both cases, the material 1 axis coincided with the x axis. All test cases started with initial condition $\mathbf{Q}(x, z, 0) = 0$, and the domain was excited by a point source with a Ricker wavelet profile having peak frequency $f_{\text{src}} = 3730$ Hz for sandstone and 3135 Hz for glass/epoxy, acting with peak intensity $+1 \text{ Pa} \cdot \text{m}^2/\text{s}$ on the vertical normal stress σ_{zz} and $-1 \text{ Pa} \cdot \text{m}^2/\text{s}$ on fluid pressure. The peak of the wavelet occurred 0.4 ms after the start of the simulation. For each case, we used a square domain 18.7 m on a side; the point source was placed at the center of the domain. The simulation time span was 1.56 ms for the sandstone medium, and 1.80 ms for glass/epoxy. We calculated results both with and without viscosity present.

We carried out all of these simulations on a uniform 501×501 cell grid. The odd number of grid cells in each axis allowed us to apply the point source to a single grid cell. A grid resolution this high was necessary to resolve the slow P wave well; as in the plane wave test cases, the fast P and S waves were well-resolved on substantially coarser grids. The point source was implemented numerically as part of the source step in the operator splitting scheme; since the point source acts on stress variables, and the viscous dissipation acts on velocity variables, it does not matter which is applied first. The CFL number for all simulations was 0.9, resulting in 279 timesteps being taken for the sandstone case and 282 for the glass/epoxy. We used the monotonized centered (MC) limiter here — even though most of the solution is smooth, without a limiter the lack of smoothness at the source point produces substantial spurious oscillations in the solution. Figure 4.5 shows the results of these simulations. These figures correspond to Figures 6 and 7 of de la Puente et al. [23], or Figures 5 and 7 of Carcione [12], and are in agreement with them.




Fig. 4.5: Results for point source test cases. Top: sandstone; bottom: glass/epoxy; left to right: v_x without viscosity, v_x with viscosity, v_z without viscosity, v_z with viscosity.

4.5. Heterogeneous domain results. We conclude our results here with a wave reflection and interconversion problem that demonstrates the ability of our code to model material interfaces, as well as the benefits of using adaptive mesh refinement — a feature of the AMRCLAW variant of CLAWPACK that we have not used thus far.

Our final test case is a large-scale inviscid problem with a bed of isotropic shale overlying isotropic sandstone, with material properties given in Table 4.1. The problem domain is the rectangle $[0, 1500 \text{ m}] \times [0, 1400 \text{ m}]$ in the x - z plane, with the boundary between the materials at $z = 700 \text{ m}$. There is a point source at $(x, z) = (750 \text{ m}, 900 \text{ m})$, again with a Ricker wavelet profile in time, with peak frequency 50 Hz. The source acts on the z -direction normal stress and fluid pressure with peak intensities $+2.3 \times 10^{13} \text{ Pa} \cdot \text{m}^2/\text{s}$ and $-2.3 \times 10^{13} \text{ Pa} \cdot \text{m}^2/\text{s}$, respectively, similarly to the previous test cases in homogeneous domains. This source magnitude was chosen to match the magnitude of the response shown by de la Puente et al. [23]. The peak of the source was delayed 40 ms after the start of the simulation, and the total duration of the run was 0.5 s. In addition to time-snapshots of the solution at particular instants, we also recorded solution time histories at three “gauges”, located at $(x_1, z_1) = (950 \text{ m}, 750 \text{ m})$, $(x_2, z_2) = (950 \text{ m}, 650 \text{ m})$, and $(x_3, z_3) = (950 \text{ m}, 500 \text{ m})$.

For our adaptively refined simulation, the coarsest-level computational grid was 75×70 cells in size, giving square cells 20 m on a side. We used two additional levels of refinement on this grid, the first at a factor of 4, and the second at further factor of 6, so that cells on the finest grids were 0.83 m on a side. Our code flagged a cell for mesh refinement when the energy norm (using the material properties of that cell) of the difference $\Delta \mathbf{Q}$ between its state vector and that of any adjacent cell exceeded $32.5 \text{ J}^{1/2}/\text{m}^{3/2}$, with the exception of the rectangle $[700 \text{ m}, 1000 \text{ m}] \times [450 \text{ m}, 950 \text{ m}]$, where the threshold for refinement was lowered to $3.25 \text{ J}^{1/2}/\text{m}^{3/2}$ in order to improve accuracy at the gauges. These tolerances were chosen empirically based on the observed magnitude of the waves in the simulation. This refinement criterion is a generalization of the typical AMRCLAW approach of refining based on the difference between the solution values in neighboring cells, which has been used successfully on many problems; AMRCLAW makes it easy to set alternate user-specified refinement criteria if desired, and also offers automatic error estimation via Richardson extrapolation. Besides refinement based on the solution field, the source location was also flagged

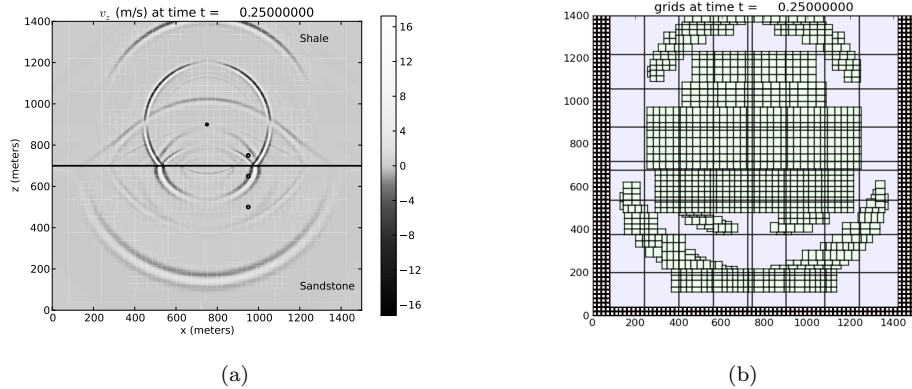


Fig. 4.6: Left: snapshot of z -direction solid velocity 0.25 seconds after the start of the simulation. The source location is marked with a solid black dot, while the gauges are marked with white-centered black dots. The gauges are numbered from top to bottom. Right: AMR grids at this time point. Individual cells are drawn on the coarsest AMR level, but only grid outlines are shown on finer levels.

for refinement to the finest level available whenever the source intensity was greater than about 10^{-9} of peak. Since the grid size was even in each direction on all but the coarsest grids, the source was distributed over the four cells closest to its location using a bilinear weighting. We again used the MC limiter for this problem.

Figure 4.6 shows a snapshot of the z -direction solid velocity 0.25 s after the start of the simulation, analogous to Figure 9(a) of de la Puente et al. [23], along with the AMR grids at this time. The solid black dot indicates the source location, and the white-centered black dots indicate the gauge locations. Because the eigenvectors associated with each wave family are different in the two materials, when a wave impinges on the material interface it produces reflected and transmitted waves in each of the three families; this results in a rich and complex solution structure. In addition to the reflected and transmitted waves, we also see head waves, where a wave in the lower half of the domain excites a slower wave family in the upper half. This results in a straight wavefront, rather than a curved one. At this point in the simulation the level 2 AMR grids have expanded to cover most of the domain, but the level 3 grids are concentrated around the wavefronts. Figure 4.7 shows the time histories of the solid x and z velocities at the three gauges. The results are in generally good agreement with Figure 10 of de la Puente et al. [23], also shown in Figure 4.7, although the peaks of slow P wave event at $t = 0.45$ s at gauge 3 are clipped in our simulation because of the limiter, and the magnitude of the second large excursion in vertical velocity at gauge 3 in our solution seems somewhat less.

While the variety of different wave speeds and reflected/transmitted waves in this problem make adaptive mesh refinement less useful than is typically the case — by halfway through the simulation time, Figure 4.6b shows a large portion of the domain refined at the finest level, because there are wavefronts present throughout the domain — we still realize a substantial savings in computation time. On an Amazon EC2 Cluster 8XL instance, running with 32 OpenMP threads, these results took 20 minutes 31 seconds to obtain, whereas a uniformly refined grid with the same cell size as the finest AMR grids took 47 minutes 19 seconds and produced no significant change in the solution. (Both times are the average of two runs; each pair differed by 4 seconds or less.) The large number of hardware threads available on this type of EC2 instance is the reason why there are so many separate fine grids in Figure 4.6b — AMRCLAW uses a coarse-grained parallelization strategy, with each grid at each time step processed by a single thread, which means that many grids must be present in order to take full advantage of highly parallel computers. The number of grids used at each refinement level is indirectly controlled by setting the maximum size of the individual grids; for the AMR computation shown above, grids were allowed to extend no more than 60 cells in any direction. Having a very large number of small grids (1047 level 3 grids in the figure) also eases load balancing between threads, since most grids are of similar size, and each thread processes many grids.

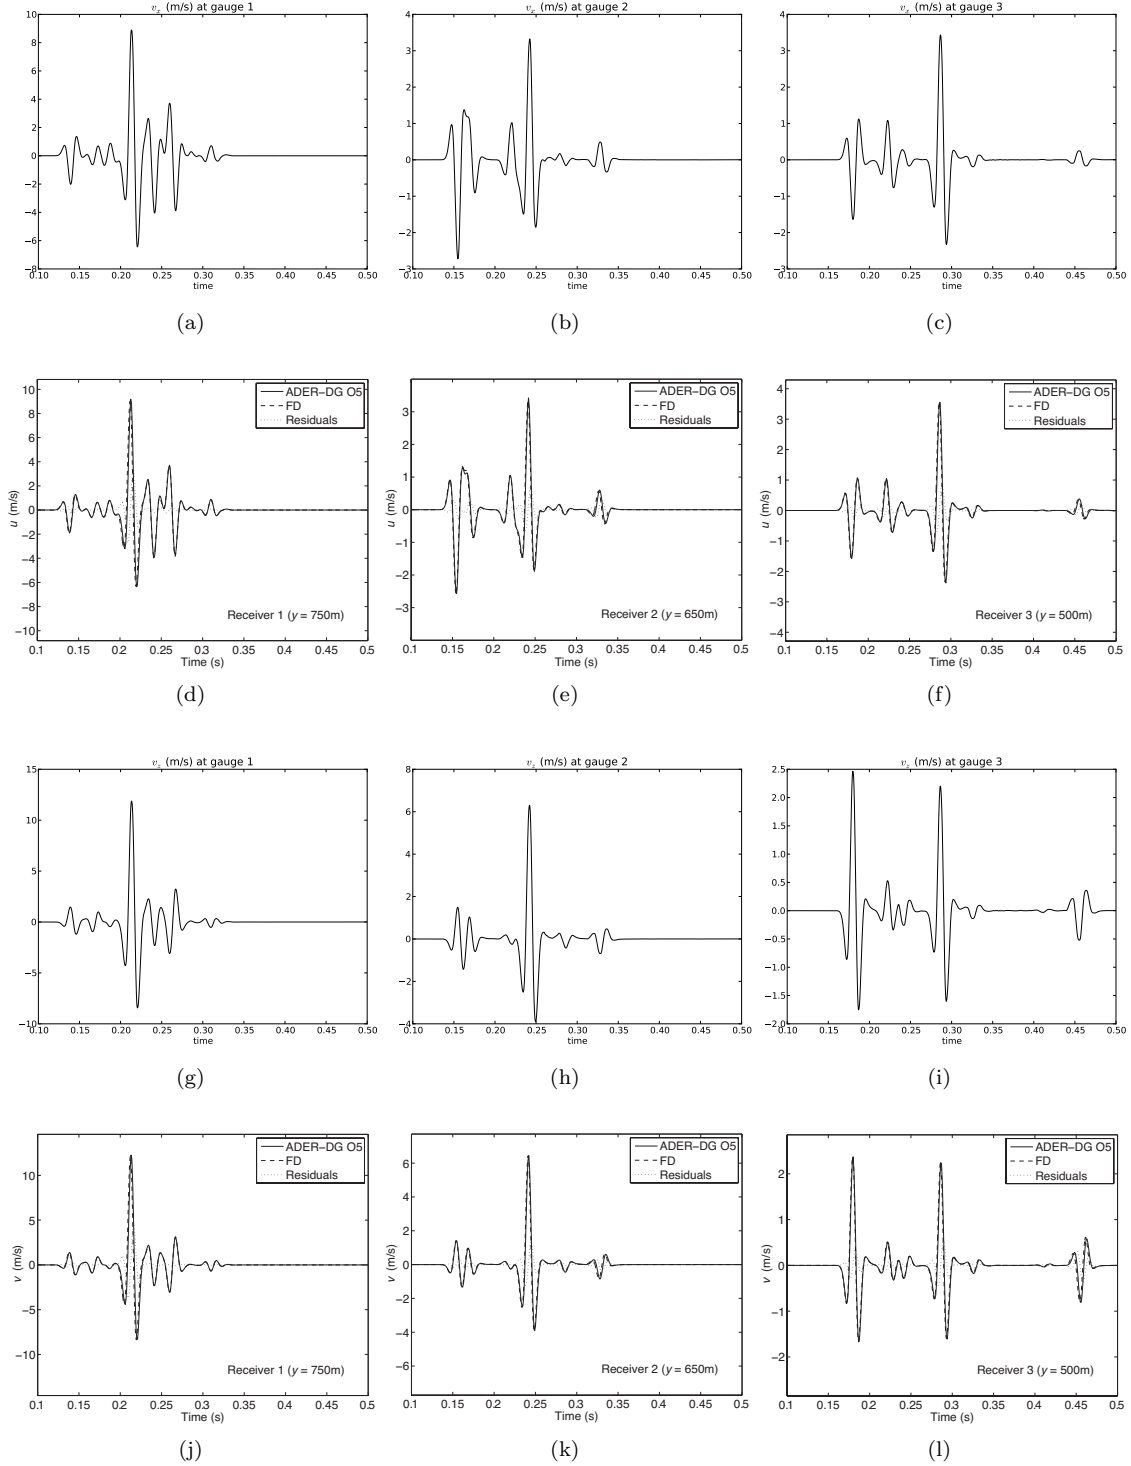


Fig. 4.7: Time-histories of solid velocity at the gauges indicated in Figure 4.6a. Top to bottom: x -direction from CLAWPACK; x -direction from de la Puente et al. [23]; z -direction from CLAWPACK; z -direction from de la Puente et al. Left to right: gauge 1, at $(x, z) = (950 \text{ m}, 750 \text{ m})$; gauge 2, at $(950 \text{ m}, 650 \text{ m})$; gauge 3, at $(950 \text{ m}, 500 \text{ m})$. Results of de la Puente et al. are reproduced in accordance with the policies of the publishing journal; the ADER-DG and FD curves in these plots are discontinuous Galerkin and finite difference results, respectively.

5. Summary and Conclusions. We have demonstrated a high-resolution finite volume code for modeling wave propagation in porous media using Biot poroelasticity theory, with the ability to model inhomogeneous domains and use adaptive mesh refinement to improve solution accuracy at substantially lower computational cost than for a uniformly refined grid. We included the dissipative source term in Biot’s equations using operator splitting. While this technique has produced spurious solutions when applied to certain types of stiff source terms in the past, we have presented an heuristic argument, based on the wave speeds of a reduced system satisfying a certain subcharacteristic condition, that we should not expect to encounter spurious solutions here. This expectation was borne out by our numerical results.

For the inviscid and viscous high-frequency regimes, our numerical solutions converge to analytic plane wave solutions for fast P and S waves with second-order accuracy. Convergence rates were somewhat impaired for the inviscid slow P wave cases tested due to the short wavelength at the frequency chosen, which caused the waves to be underresolved on the coarsest grids. However, the viscous, high-frequency slow P wave test cases showed unambiguous second-order convergence, which indicates that we can also expect second order in the inviscid case when the slow P wave is well-resolved. Due to the relative stiffness of the source term compared to the problem timescale for low-frequency waves, we obtained only roughly first-order accuracy for fast P and S waves in the low-frequency viscous regime. We obtained second-order accuracy for low-frequency slow P waves, but the slow P test cases were not directly comparable to the fast P and S cases. The other test cases examined, involving results for a point source in either a homogeneous orthotropic medium, or in a layered bed of two distinct isotropic media, agreed well with results for the same test cases published by other authors.

There are substantial opportunities for future work based on what we have presented here. One possibility is the extension of the finite volume methods used here to logically rectangular mapped grids. This extension is straightforward, and allows modeling of more complex domains with internal boundaries, so long as a (not necessarily smooth) mapping function can be found that maps the internal and external boundaries to rectangles in the computational domain. Another opportunity for future work is the implementation of a more accurate solution procedure at low frequencies, such as one based on the methods discussed by Hittinger [34] or Pember [42]. We intend to explore both these avenues in subsequent publications, along with applying the software developed here to some specific problems.

To aid in the reproducibility of the results presented here, we provide all of the code used to generate them at https://bitbucket.org/grady_lemoine/poro-2d-cartesian-archive.

REFERENCES

- [1] E. ANDERSON, Z. BAI, C. BISCHOF, S. BLACKFORD, J. DEMMEL, J. DONGARRA, J. DU CROZ, A. GREENBAUM, S. HAMMARLING, A. MCKENNEY, AND D. SORENSEN, *LAPACK Users’ Guide*, Society for Industrial and Applied Mathematics, Philadelphia, PA, third ed., 1999.
- [2] KEITH ATTENBOROUGH, DAVID L. BERRY, AND YU CHEN, *Acoustic scattering by near-surface inhomogeneities in porous media.*, tech. report, Defense Technical Information Center OAI-PMH Repository [<http://stinet.dtic.mil/oai/oai>] (United States), 1998.
- [3] M. J. BERGER AND R. J. LEVEQUE, *Adaptive mesh refinement using wave-propagation algorithms for hyperbolic systems*, SIAM Journal on Numerical Analysis, 35 (1998), pp. 2298–2316.
- [4] M. A. BIOT, *Theory of propagation of elastic waves in a fluid-saturated porous solid. I. Low-frequency range*, Journal of the Acoustical Society of America, 28 (1956), pp. 168–178.
- [5] ———, *Theory of propagation of elastic waves in a fluid-saturated porous solid. II. Higher frequency range*, Journal of the Acoustical Society of America, 28 (1956), pp. 179–191.
- [6] ———, *Mechanics of deformation and acoustic propagation in porous media*, Journal of Applied Physics, 33 (1962), pp. 1482–1498.
- [7] JAMES L. BUCHANAN AND ROBERT P. GILBERT, *Determination of the parameters of cancellous bone using high frequency acoustic measurements*, Mathematical and Computer Modelling, 45 (2007), pp. 281–308.
- [8] ———, *Determination of the parameters of cancellous bone using high frequency acoustic measurements II: inverse problems*, Journal of Computational Acoustics, 15 (2007), pp. 199–220.
- [9] JAMES L. BUCHANAN, ROBERT P. GILBERT, AND KHALDOUN KHASHANAH, *Determination of the parameters of cancellous bone using low frequency acoustic measurements*, Journal of Computational Acoustics, 12 (2004), pp. 99–126.
- [10] J. L. BUCHANAN, R. P. GILBERT, A. WIRGIN, AND Y. S. XU, *Marine acoustics: direct and inverse problems*, SIAM, Philadelphia, 2004.
- [11] M. ČADA AND M. TORRILHON, *Compact third-order limiter functions for finite volume methods*, Journal of Computational Physics, 228 (2009), pp. 4118–4145.

- [12] J. M. CARCIONE, *Wave propagation in anisotropic, saturated porous media: plane-wave theory and numerical simulation*, Journal of the Acoustical Society of America, 99 (1996), pp. 2655–2666.
- [13] J. M. CARCIONE, *Wave Fields in Real Media: Wave Propagation in Anisotropic, Anelastic, and Porous Media*, Elsevier, Oxford, 2001.
- [14] J. M. CARCIONE, C. MORENCY, AND J. E. SANTOS, *Computational poroelasticity – a review*, Geophysics, 75 (2010), pp. 75A229–75A243.
- [15] J. M. CARCIONE AND G. QUIROGA-GOODE, *Some aspects of the physics and numerical modeling of Biot compressional waves*, Journal of Computational Acoustics, 3 (1995), pp. 261–280.
- [16] G.-Q. CHEN, C. D. LEVERMORE, AND T.-P. LIU, *Hyperbolic conservation laws with stiff relaxation terms and entropy*, Communications in Pure and Applied Mathematics, 47 (1994), pp. 787–830.
- [17] G. CHIAVASSA AND B. LOMBARD, *Time domain numerical modeling of wave propagation in 2D heterogeneous porous media*, Journal of Computational Physics, 230 (2011), pp. 5288–5309.
- [18] P. COLELLA, A. MAJDA, AND V. ROYTBURD, *Theoretical and numerical structure for reacting shock waves*, SIAM Journal on Scientific and Statistical Computing, 7 (1986), pp. 1059–1080.
- [19] S. C. COWIN, *Bone poroelasticity*, Journal of Biomechanics, 32 (1999), pp. 217–238.
- [20] S. C. COWIN AND L. CARDOSO, *Fabric dependence of bone ultrasound*, Acta of Bioengineering and Biomechanics, 12 (2010).
- [21] S. C. COWIN AND M. M. MEHRABADI, *Identification of the elastic symmetry of bone and other materials*, Journal of Biomechanics, 22 (1989), pp. 503–515.
- [22] N. DAI, A. VAFIDIS, AND E. KANASEWICH, *Wave propagation in heterogeneous porous media: a velocity-stress, finite-difference method*, Geophysics, 60 (1995), pp. 327–340.
- [23] J. DE LA PUENTE, M. DUMBSER, M. KÄSER, AND H. IGEL, *Discontinuous Galerkin methods for wave propagation in poroelastic media*, Geophysics, 73 (2008), pp. T77–T97.
- [24] G. DEGRANDE AND G. DE ROECK, *FFT-based spectral analysis methodology for one-dimensional wave propagation in poroelastic media*, Transport in Porous Media, 9 (1992), pp. 85–97.
- [25] H. DERESIEWICZ AND R. SKALAK, *On uniqueness in dynamic poroelasticity*, Bulletin of the Seismological Society of America, 53 (1963), pp. 783–788.
- [26] E. DETOURNAY AND A. H.-D. CHENG, *Poroelastic response of a borehole in a non-hydrostatic stress field*, International Journal of Rock Mechanics and Mining Sciences and Geomechanics Abstracts, 25 (1988), pp. 171–182.
- [27] S. K. GARG, A. H. NAYFEH, AND A. J. GOOD, *Compressional waves in fluid-saturated elastic porous media*, Journal of Applied Physics, 45 (1974), pp. 1968–1974.
- [28] R. F. GIBSON, *Principles of Composite Material Mechanics*, McGraw-Hill, New York, 1994.
- [29] R. P. GILBERT, P. GUYENNE, AND M. YVONNE OU, *A quantitative ultrasound model of the bone with blood as the interstitial fluid*, Mathematical and Computer Modelling, 55 (2012), pp. 2029–2039.
- [30] R. P. GILBERT AND Z. LIN, *Acoustic field in a shallow, stratified ocean with a poro-elastic seabed*, Zeitschrift für Angewandte Mathematik und Mechanik, 77 (1997), pp. 677–688.
- [31] R. P. GILBERT AND M. YVONNE OU, *Acoustic wave propagation in a composite of two different poroelastic materials with a very rough periodic interface: a homogenization approach*, International Journal for Multiscale Computational Engineering, 1 (2003).
- [32] BORIS GUREVICH AND MICHAEL SCHOENBERG, *Interface conditions for Biot’s equations of poroelasticity*, Journal of the Acoustical Society of America, 105 (1999), pp. 2585–2589.
- [33] S. HASSANZADEH, *Acoustic modeling in fluid-saturated porous media*, Geophysics, 56 (1991), pp. 424–435.
- [34] J. A. HITTINGER, *Foundations for the generalization of the Godunov method to hyperbolic systems with stiff relaxation source terms*, PhD thesis, University of Michigan, 2000.
- [35] F. KEMM, *A comparative study of TVD-limiters – well-known limiters and an introduction of new ones*, International Journal for Numerical Methods in Fluids, 67 (2011), pp. 404–440.
- [36] R. J. LEVEQUE, *Finite Volume Methods for Hyperbolic Problems*, Cambridge University Press, New York, 2002.
- [37] R. J. LEVEQUE AND H. C. YEE, *A study of numerical methods for hyperbolic conservation laws with stiff source terms*, Journal of Computational Physics, 86 (1990), pp. 187–210.
- [38] X.-D. LIU AND E. TADMOR, *Third order nonoscillatory central scheme for hyperbolic conservation laws*, Numerische Mathematik, 79 (1998), pp. 397–425.
- [39] B. G. MIKHAILENKO, *Numerical experiment in seismic investigations*, Journal of Geophysics, 58 (1985), pp. 101–124.
- [40] C. MORENCY AND J. TROMP, *Spectral-element simulations of wave propagation in porous media*, Geophysical Journal International, 179 (2008), pp. 1148–1168.
- [41] R. B. PEMBER, *Numerical methods for hyperbolic conservation laws with stiff relaxation I. Spurious solutions*, SIAM Journal on Applied Mathematics, 53 (1993), pp. 1293–1330.
- [42] ———, *Numerical methods for hyperbolic conservation laws with stiff relaxation II. Higher-order Godunov methods*, SIAM Journal on Scientific Computing, 14 (1993), pp. 824–859.
- [43] J. E. SANTOS AND E. J. OREÑA, *Elastic wave propagation in fluid-saturate porous media, part II: The Galerkin procedures*, Mathematical Modeling and Numerical Analysis, 20 (1986), pp. 129–139.
- [44] M. D. SHARMA, *Wave propagation across the boundary between two dissimilar poroelastic solids*, Journal of Sound and Vibration, 314 (2008), pp. 657–671.
- [45] THE CLAWPACK AUTHORS, *CLAWPACK software*. www.clawpack.org.